AtoMix: Fostering Structured Visual Ideation for Remote Groups through Atomic Composition and Cross-Pollination

Heeji Jolie Kim
Industrial Design
KAIST
Daejeon, Republic of Korea
jolieheejikim@kaist.ac.kr

Anam Ahmad Khan
Industrial Design
KAIST
Daejeon, Republic of Korea
anam.khan@kaist.ac.kr

SeungHui Huh
Industrial Design
KAIST
Daejeon, Republic of Korea
seunghui@kaist.ac.kr

Andrea Bianchi
Industrial Design
KAIST
Daejeon, Republic of Korea
andrea@kaist.ac.kr

Abstract

6-3-5 brainwriting is a structured, silent ideation method where 6 participants each generate 3 ideas in 5 minutes across 6 rounds, yielding up to 108 ideas and enabling equal participation, rapid ideation and cumulative development. Many ideas are fundamentally visual, making generative image models especially useful for accelerating visual exploration and making abstract ideas quickly tangible, but existing tools cannot fully support visual composition, where ideas are built from atomic elements and iteratively remixed. A formative study (N=18) showed that even with GenAI embedded in a 6-3-5 structure, participants invested effort in crafting monolithic prompts yet struggled to obtain intended results, build on others' ideas, or work modularly, leading to redundancy and limited divergence. To address these issues, we present AtoMix, a collaborative visual interface for remote 6-3-5 brainwriting that supports granularity and cross-pollination through atomic prompting, canvas composition, and reuse of image fragments. In a comparative study (N=24), AtoMix afforded more fine-grained interaction than the baseline and better supported collaborative visual ideation.

CCS Concepts

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

Keywords

Creativity Support Tools (CST), Visual Ideation, Brainwriting, Generative Image

ACM Reference Format:
Heeji Jolie Kim, Anam Ahmad Khan, SeungHui Huh, and Andrea Bianchi. 2026. AtoMix: Fostering Structured Visual Ideation for Remote Groups through Atomic Composition and Cross-Pollination . In *Creativity and Cognition (C&C '26), July 13–16, 2026, London, United Kingdom.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3803784.3807530

ACM ISBN 979-8-4007-2583-8/2026/07
https://doi.org/10.1145/3803784.3807530

1 Introduction

Ideation is central to creative work in design. Established frameworks such as brainstorming [50], brainwriting [14], and brainsketching [54] structure how ideas emerge in early innovation. Brainstorming, while popular, suffers from well-documented drawbacks: evaluation apprehension, production blocking, social loafing, and illusions of productivity all constrain the number of unique ideas relative to nominal (solo) ideation [14]. Brainwriting was introduced to address these issues by having participants silently write and exchange ideas, reducing status and dominance effects which supports more reflective engagement with others' contributions [14]. Prior work has extended brainwriting into visual domains: brainsketching uses hand-drawn sketches, rather than speech or text, as the primary medium of thought, with participants drawing individually, rotating positions, and iteratively building on one another's sketches [54]. Compared to text-based brainwriting, visual approaches provoke new ideas through cycles of 'seeing-making-seeing', described as a "reflective conversation" [40] with visual materials [54]. We use the term *visual brainwriting* to describe the process of conducting brainwriting through incorporating visual media instead of merely using written text.

Recent work shows that *AI-augmented brainwriting* can further strengthen early ideation by increasing idea fluency, accelerating elaboration and reframing, and supporting movement between divergence and convergence [41]. As ideas grow more complex, AI-augmented ideation can begin to break down: generated suggestions often drift from user intent, requiring repeated regeneration and repair [8, 48]. Moreover, while AI can increase idea fluency, it may also promote early convergence by steering participants toward a narrow set of coherent ideas [15, 43]. These tensions raise questions about how AI should be integrated into ideation practices without undermining their core strengths. Motivated by these findings, we set out to examine whether these limitations of AI-augmented brainwriting also arise in Generative AI (GenAI), specifically image generation, supported *visual brainwriting*, where ideas are developed through iterative interaction with visual materials rather than text alone. To explore this, we first conducted a formative study (N=18) using a digital visual brainwriting board following a 6–3–5 framework augmented with generative image tools to observe how particpants use GenAI with visual brainwriting task. Our results suggests that participants fell into prompt–regenerate loops that treated AI outputs as finished artifacts rather than manipulable materials, and AI-generated images fell short when ideas required multi-component or process-level reasoning. Together, these frictions shifted ideation toward whole-image regeneration and convergence, undermining visual brainwriting's core strengths of incremental iteration, externalization, and cumulative build-on.

To address these limitations and better realize visual brainwriting's potential in digital environments, we propose a granular, fragment-level approach grounded in prior evidence that ambiguous, manipulable visual materials support iterative reinterpretation, idea discovery, and incremental build-on connections [51, 53, 54]. We introduce AtoMix, a GenAI-augmented visual brainwriting system that re-centers *composition over completion.* AtoMix provides (i) *collage-based canvas* to generate object-level visuals that drop into the canvas as movable layers; (ii) *region-based editing* that aligns new content to user-selected areas without replacing the whole; and (iii) *cross-canvas mashups* that surface recombination opportunities and carry provenance for traceable build-on. These affordances may help mitigate the drift toward whole-image regeneration and early convergence by enabling atomic, incremental, and cumulative visual ideation, while still leveraging AI generation for speed and fluency.

To explore how these affordances shape visual brainwriting practice, we conducted a comparative study (N=24) examining artifact characteristics and interaction patterns in AtoMix versus a baseline system. Our study showed that AtoMix produced canvases with greater compositional richness, higher vertical continuity across rounds, stronger horizontal differentiation between ideas, and fewer empty canvases compared to the baseline. These outcomes were driven by three emergent interaction patterns: atomic prompting (shorter, more frequent prompts), granular editing (region-based modifications), and cross-pollination (reusing and remixing fragments across canvases). Overall, this research makes the following contributions:

- AtoMix, a collaborative visual brainwriting system enabling atomic composition, granular editing, and cross-canvas recombination;
- a design approaches that positions GenAI outputs as compositional materials rather than complete images, supporting iterative assembly over monolithic generation;
- empirical findings from a comparative study (N=24) examining output characteristics (compositional expressiveness, iterative development, idea differentiation), interaction patterns (atomic prompting, granular editing, cross-pollination), and user experience in AtoMix versus baseline conditions.

2 Background and Related Works

We aim to propose a workflow for visual brainwriting that helps designers work more efficiently in group-based ideation practices, using generative AI to enable and accelerate the visualization of ideas. To ground this approach, we review three strands of prior work: (1) *brainwriting and visual brainwriting* as structured, group-oriented ideation methods, (2) *creativity support tools (CSTs)* that scaffold ideation, and (3) *GenAI approaches* that augment ideation, including visualization of ideas and the generation/recombination of concepts.

2.1 Brainwriting

Brainwriting addresses traditional brainstorming challenges— production blocking, evaluation apprehension, and social loafing—by enabling participants to generate and exchange ideas in parallel rather than speaking sequentially [3, 14, 28]. A common framework is 6–3–5, where six participants each produce three ideas and pass their sheets every five minutes, ensuring structured rotation and equal contribution [36, 42]. *Visual brainwriting* extends this approach by using visual media as the unit of thought. The original form, brainsketching [54], has participants sketch individually, then rotate to build on others' marks, creating incremental "seeing–making–seeing" cycles that promote reinterpretation and cumulative development [51, 53, 54]. These cycles yield partial, evolving artifacts rather than finished drawings, supporting divergent exploration and stepwise refinement, which are properties we seek to preserve in visual brainwriting. Central to this process is *cross-pollination*: the reuse and recombination of ideas across participants, a form of cross-fertilization where exposure to others' contributions stimulates new associations [26, 34]. In brainwriting, cross-pollination occurs structurally through sheet rotation, enabling participants to build on, reinterpret, or remix prior work. These dynamics support *divergent thinking*—the generation of multiple, varied solutions from a single starting point [12, 37]—alongside *combinational creativity*, where novel ideas emerge from blending existing concepts [2]. The strict timing further constrains elaborate rendering, encouraging quick, rough visuals that leave interpretive space for others and transform initially separate sheets into an increasingly interconnected pool of ideas [51, 53, 54]. However, visual brainwriting also surfaces a persistent *visualization gap*: participants often struggle to express intended concepts visually. They may hesitate to draw in front of others ("I can't draw"), while time-bounded rounds heighten performance pressure and complex, multi-component ideas resist being compressed into a single image [11, 42, 54].

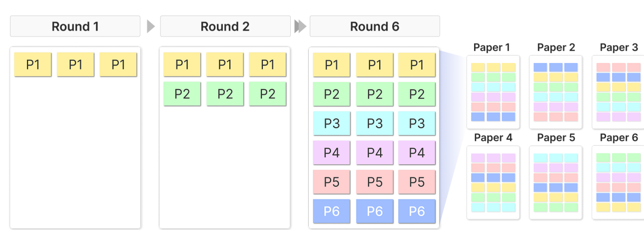

Figure 1: Illustration of the 6–3–5 visual brainwriting method. Six participants (P1–P6) work in parallel; in each 5-minute round, every person creates *three* visual ideas on a sheet and then passes the sheet to the next participant. As rounds progress (Round 1 → Round 6), each sheet accumulates layered contributions from all participants. The outcome is six papers, each showing a grid of ideas that reflect incremental build-on [42, 54].

This visualization challenge highlights the potential for creativity support tools to lower these barriers while preserving visual brainwriting's key benefits: incremental idea development, collaborative building across participants, and the interpretive flexibility of rough, partial visuals. Our work aims to address the visualization gap through the design of a creativity support tool for visual brainwriting.

2.2 CSTs for Ideation

Creativity support tools (CSTs) are interactive systems designed to enhance human creativity across a range of activities, with one important strand focusing on scaffolding ideation through exploration, expression, and problem-solving [10, 45]. Shneiderman's framework identifies core principles for effective creativity support—low thresholds, wide walls, and high ceilings—alongside the importance of exploration, collaboration, and history-keeping [45]. Many CSTs also build on *direct manipulation* [44], where continuous, reversible operations on visible objects reduce cognitive overhead and invite experimentation. Below, we contrast two strands of CSTs in ideation—*individual* and *collaborative* tools—to identify gaps our work addresses.

2.2.1 Individual Ideation Support. CSTs for individual ideation embed context-sensitive cues, domain-specific examples [6, 18], and mechanisms for *curated surprise*: graph-based tools expose non-obvious links [1], pen+touch diagramming maintains productive ambiguity [58], and semantic retrieval steers users toward diverse ideas [46]. However, most systems operate at the level of *complete* examples or query results rather than manipulable parts, offering limited support for modular recombination or cumulative build-on [51, 53, 54]. As a result, how individually expanded idea spaces translate into shared, incremental ideation remains underexplored [10, 45].

2.2.2 Collaborative Ideation Support. Collaborative CSTs strengthen group brainstorming through coordination, visibility, and real-time artifact manipulation. Tangible approaches increase social presence during remote sessions [33] and structure turn-taking [32]. Sketch-centric tools support co-located brainstorming [21], moderated remote sketching [29], and hybrid workspaces linking paper with interactive layers [11]. Group-oriented platforms provide shared canvases with branching and merging [20, 62], structured ideation grids [24], game-based interfaces [16], and semantic modeling [46]. These systems suggest that collaboration improves when participants can interact rapidly and incrementally with partial artifacts that others can see and extend.

Takeaways on CSTs for Ideation. Individual CSTs expand personal idea spaces; collaborative platforms make thinking visible in shared workspaces [10, 45]. Despite the availability of direct manipulation [44] and creativity support principles [45], most tools still treat ideas as complete products rather than *manipulable materials* that can be broken apart, modified, and reused—the very mechanics that sustain cumulative build-on in visual brainwriting [51, 53, 54].

2.3 GenAI for Early-Stage Ideation

GenAI has lowered barriers to visual exploration through text-to-image systems like ChatGPT[1], Stable Diffusion[2], and Midjourney[3], enabling rapid generation of intended or unexpected visuals that stimulate divergent thinking [36, 39].

2.3.1 GenAI for Idea Generation. GenAI-for-ideation splits into two trajectories. One embeds AI into *group formats*, extending brainwriting with models that propose, cluster, or evaluate ideas [41, 61]. Another targets *designer control* through iterative exploration, reference reuse, and prompt scaffolds [27, 57]. Adjacent work positions AI as a *sketching partner*: co-creative tools return complementary sketches using analogy to widen the search space [22, 30, 60], while domain tools leverage mood boards and concept blending [7, 17, 49, 55]. However, these systems typically treat ideas as *whole renders* or *keyword bundles*, with recombination remaining a single-user activity [14, 51, 53]. This monolithic interaction encourages regenerate-and-replace workflows, limiting shared buildup across people and rounds.

2.3.2 GenAI for Visualization. Generative models increasingly turn abstract concepts into visual forms. Within this space, one strand streamlines prompt interaction through tools that decompose and blend concepts or expand the design space via prompt suggestion [5, 47, 56], while another emphasizes recombination, blending image elements or integrating generation with reference retrieval [7, 49, 57]. Parallel to these research tools, modular generative workflows such as ComfyUI[4] expose generation pipelines as node graphs, giving users fine-grained control over individual processing steps; however, these environments target technical users and do not support real-time collaborative ideation. Building on these directions, collaborative tools further combine mood boards with semantic search and biological analogies [13, 19, 25, 27].

Takeaways on GenAI for Ideation. What remains under- explored is not *whether* AI can produce visuals, but *how* those visuals function as manipulable materials in collaborative ideation. Current systems return *complete renders*—effective for seeding ideas but inadequate for visual brainwriting's mechanics: (i) incremental composition of partial visuals rather than finished images; (ii) situated extension through in-place transformation rather than regeneration; and (iii) idea circulation where visible lineage sustains reinterpretation across participants [51, 53, 54]. Current GenAI is optimized for *finished images*, whereas visual brainwriting gains power from *manipulating elements* across time and collaborators.

3 Formative Study

To examine the challenges and opportunities of visual brainwriting with generative images, we conducted a formative study using a digital adaptation of the 6–3–5 brainwriting method [59]. This adaptation allowed us to observe how participants used visual content and generative tools during sequential ideation, informing the design of systems for multi-round idea development.

[1]https://chatgpt.com/
[2]https://stablediffusionweb.com/
[3]https://www.midjourney.com/
[4]https://www.comfy.org/

3.1 Participants and Procedure

We recruited 18 participants (9 female, 9 male), aged 20–27 (M = 22.4, SD = 2.52), from the institutional community via online postings. Participants represented diverse academic backgrounds, including industrial design, electrical engineering, computer science, and chemistry.

The brainwriting activity was conducted in a Figma[5] workspace replicating the 6–3–5 method and adapted for collaborative visual ideation with generative AI. As shown in Fig. 2, the workspace comprised six pages, each with six rows and three columns. To support deeper exploration, we extended each round from 5 to 10 minutes. Participants could use any text or image generative AI tools, including ChatGPT[6] , Gemini[7] , and Copilot[8]

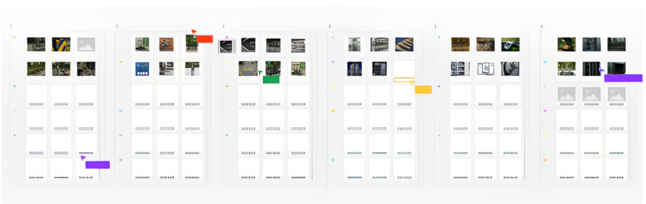

Figure 2: Figma workspace used for formative study— 6 papers with 6 rows (rounds), 3 columns (ideas), currently in round 2

Figure 3 shows the study procedure. After a brief introduction, participants accessed the Figma workspace and completed the same task: *Design a storage or parking solution for personal micro-mobility devices integrated into urban residential environments.* Using images and text, participants followed the 6–3–5 structure, passing ideas to the next participant for further development. We conducted three 2-hour sessions with six participants each. Each session included six 10-minute rounds, followed by a 40-minute semi-structured group interview exploring participants' experiences, use of AI tools, and strategies for generating and evolving visual ideas. Interviews were audio recorded, transcribed, and analyzed using thematic analysis [4].

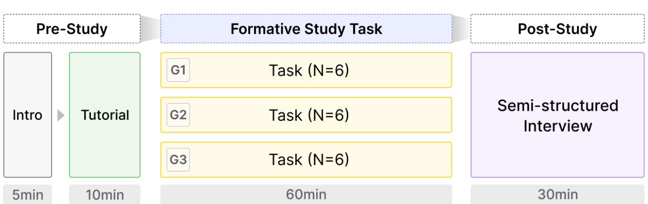

Figure 3: Formative study procedure. Each session comprised an 15-minute introduction (briefing, consent, and tutorial), a 60-minute ideation activity comprised of 6 rounds of 10 minutes, and a 40-minute semi-structured interview on participants' experiences.

[5]https://www.figma.com/
[6]https://chatgpt.com/
[7]https://gemini.google.com/
[8]https://copilot.microsoft.com/

3.2 Findings

Our analysis showed that GenAI both supported and constrained creativity in visual brainwriting. While participants valued its speed and generative capacity, they encountered intent–output mismatches, limited control, and difficulty expressing complex ideas with single images. Participants adapted by reusing, recombining, and diverging ideas, exposing gaps in existing tools and informing our design goals for balancing AI assistance with creative agency.

3.2.1 ***Modular and Collage-Based Expression.*** Participants valued AI's ability to rapidly visualize product ideas, making them easy to grasp at a glance. For instance, P8 noted that tangible concepts like a foldable kickboard were "instantly understandable." However, AI-generated images were less effective for multi-component or process-oriented ideas (P10, P11). As P10 explained, complex systems require step-by-step representation, and "trying to explain the whole process with just a single picture" often caused confusion. To compensate, participants adopted collage-like strategies, reusing prior outputs as building blocks (P11) or manually recombining images via screenshots in Figma (P12). Many also wanted finer editing control, such as cropping or enlarging specific regions to highlight key elements (P9, P12, P17). Together, these findings show that when single images were insufficient, participants decomposed, reused, and recombined visual elements, motivating the need for compositional expressiveness that support modular editing, layering, and incremental refinement for communicating complex ideas.

3.2.2 ***Limits of Control and Contextual Refinement.*** Participants valued generative AI for sparking unexpected ideas, especially in early rounds. As P8 noted, "It gave me things I wouldn't have thought of," highlighting AI's role in expanding creative possibilities. However, outputs often misaligned with user intent: refinements produced near-identical images despite prompt changes (P8, P10, P12) or introduced missing or irrelevant elements (P9, P12). Participants tried workarounds such as re-uploading images, simplifying prompts, or wishing for low-fidelity previews, which underscored the lack of fine-grained control. As P17 explained, it would be "much [more] efficient to make small manual edits rather than regenerate images from scratch". Overall, these findings show that while AI supported inspiration and variation, limited control hindered intent alignment, motivating the need for localized refinement for selective modification.

3.2.3 ***Creativity Relies on Reuse and Divergence.*** Participants sustained creativity through two complementary strategies: reusing earlier ideas and seeking divergent stimuli when momentum slowed. Reuse involved revisiting prior prompts or images, extracting useful elements, and recombining them to extend ideas without starting over (P2, P14–16). For example, P2 compiled prompts to identify "common elements and improvement points," while P14 remixed earlier images with new ones. However, participants lacked modular control; as P18 noted, "I wanted to just grab one part of an image, but there's no easy way to do that," pointing to needs for cropping, style extraction (P11), preset backgrounds, or shared libraries (P12). When ideation slowed, participants adopted divergent strategies to regain momentum, such as revisiting earlier canvases (P13, P16), feeding accumulated text into ChatGPT (P17), or switching AI tools when outputs became repetitive (P13, P17). Together, these findings show that reuse supported continuity while divergence enabled novelty, but the lack of fragment-level editing and structured recombination forced ad hoc workarounds, motivating systems that better support modular reuse, flexible recombination, and divergent exploration across rounds.

3.3 Design Goals

Based on insights from our formative study, we identified three design goals to guide the development of our collaborative brainwriting system.

- **DG1: Compositional Expression.** The system should allow users to express complex, multi-component ideas by composing visual elements modularly—selecting, layering, and reusing fragments across rounds—rather than relying on single monolithic images.
- **DG2: Localized Refinement.** The system should support fine-grained modification of specific regions or elements, preserving surrounding context and creative intent without requiring full regeneration.
- **DG3: Recombinational Flow.** The system should sustain ideation momentum by enabling recombination of visual and textual fragments across canvases, supporting both reuse and divergence.

4 AtoMix

Guided by the derived design goals, we implemented AtoMix, a web-based collaborative ideation system that integrates image generation, collage composition, granular editing, and cross-canvas mashup across multiple canvases. We first present a system overview, followed by a description of each feature aligned with the three design goals.

4.1 System Workflow

In AtoMix, we used the 6–3–5 method where participants work through six rounds, with three canvases unlocked per round for five minutes before rotating to the next person's work. Within this structure, users visualize ideas by composing generated images, making granular changes, and synthesizing ideas from multiple canvases.

4.1.1 ***Collage-Based Canvas.*** Each canvas functions as a freeform collage space where users generate images by typing descriptions into the prompt panel. Generated images appear in a queue; users drag them onto the canvas and manipulate them as independent fragments—resizing, rotating, repositioning, or deleting. The toolbar provides text boxes, a pen tool, and background removal to support layering images without bounding-box overlap. Each element is stacked on a separate layer with individual visibility control, enabling users to compose multiple images, annotations, and sketches into unified concepts (Figure 5a).

4.1.2 ***Region-Based Editing.*** Region-based editing enables fine-grained modifications without regenerating entire images. Users select the lasso tool, trace a region on an existing image, and enter a short instruction—optionally attaching a reference image. The system sends the selected patch with its local context to the generative model; the returned result is inserted as a new fragment layer aligned to the selected region, leaving the original untouched. A layer panel lets users reorder, hide, or duplicate fragments to compare alternatives (Figure 5b).

4.1.3 ***Cross-Canvas Mashups.*** By clicking on the "lightbulb" button, users can select two to four canvases and request a mashup. The system extracts snapshots and metadata, blending elements from each source. The generative model produces a new image with a textual description of the combined concept, which users insert onto their current canvas (Figure 5c).

Across all features, AtoMix emphasizes lightweight, compositional workflows that align with visual brainwriting's ethos of "thinking while doing" [54].

4.2 Implementation

AtoMix is a multi-user web application built with JavaScript (ES6) and Fabric.js[9], which provides the object model for the canvas. Each image, sketch, and annotation is treated as an independent object that can be directly manipulated rather than flattened into a bitmap, with metadata storing its originating canvas, creation time, and source prompt.

Firebase Firestore[10] maintains the state of papers, canvases, and fragments across participants, with updates synchronized in real time using incremental patches. Lineage metadata is recorded to trace origin across rounds. Generated images are stored on Cloudinary[11], with URLs stored in Firestore for retrieval in subsequent rounds.

Image generation uses the Gemini 2.5 Flash API[12]. The system wraps user prompts with uniform settings (resolution, aspect ratio, style) and augments requests with contextual metadata—for region-based edits, this includes the selected patch and surrounding context; for cross-canvas mashup, rasterized snapshots and metadata from multiple canvases are merged to guide synthesis.

5 User Study

The purpose of this study is to explore how our proposed visual brainwriting system influences idea generation and development, and whether its design features lead participants to engage in ideation differently. We conducted a between-subjects study with 24 participants, split into two groups using AtoMix and two groups using a baseline system. We analyzed emergent ideation patterns during system use and assessed participants' overall user experience.

- **RQ1**: How do ideation artifacts differ between AtoMix and the baseline system?
- **RQ2**: How do interaction patterns in AtoMix and the baseline system influence ideation artifacts?
- **RQ3**: How do participants perceive and experience AtoMix for visual brainwriting?

[9]http://fabricjs.com/
[10]https://firebase.google.com/docs/firestore
[11]https://cloudinary.com/home
[12]https://ai.google.dev/gemini-api/docs

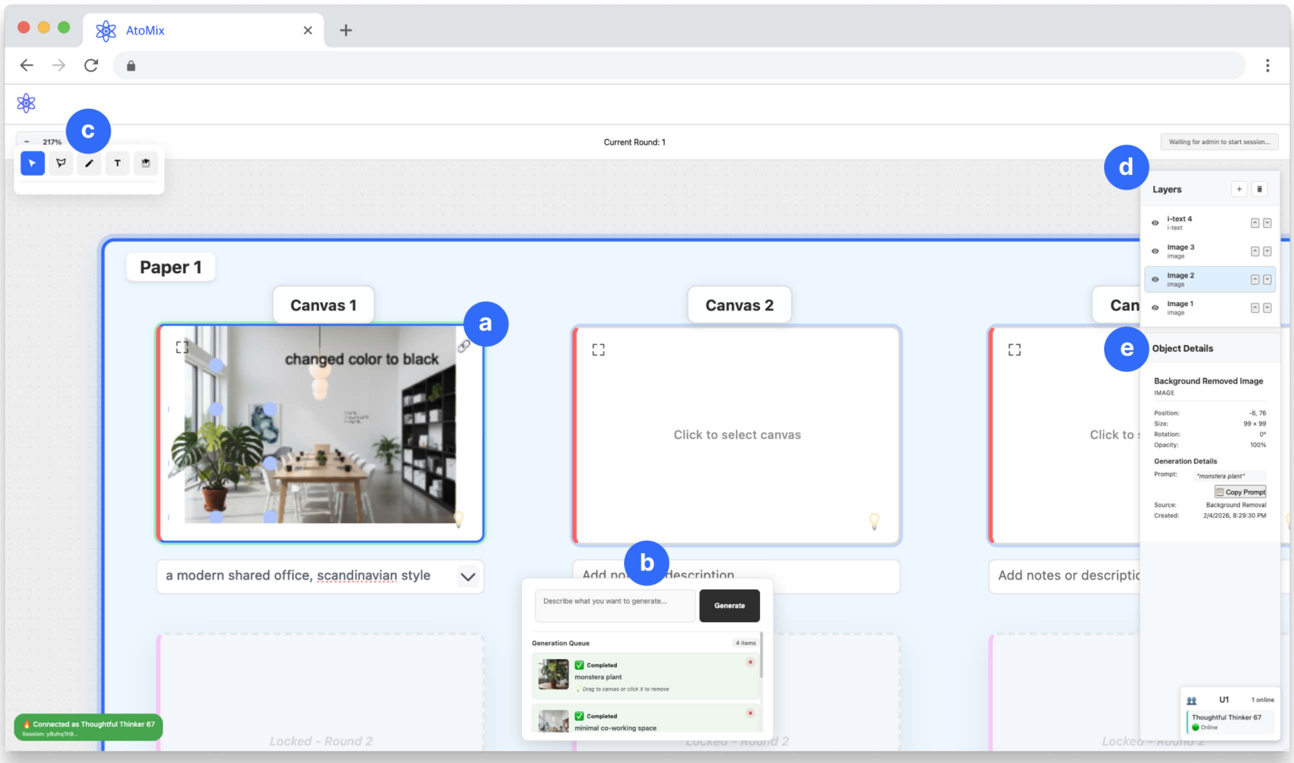

Figure 4: System Screenshot (a) Multi-canvas collage board: three side-by-side canvases per round; users arrange fragments and add per-canvas notes. (b) Prompt & Generation Queue: enter a description to generate images; completed results appear with status and can be dragged onto any canvas. (c) Toolbar: selection, lasso for region-based edits, background removal, pen, and text. (d) Layers panel: every fragment is a layer; users reorder, hide, and select layers for precise control. (e) Object Details: per-fragment metadata (prompt, source, timestamp) for provenance and copy-back of prompts. The example depicts a Scandinavian design for a shared office, featuring a plant image assembled from generated fragments and annotations.

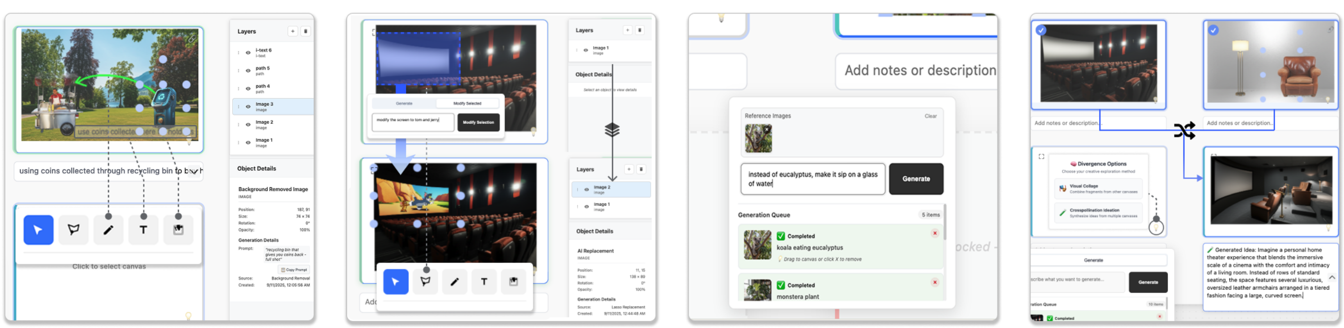

(a) Collage-Based Canvas (b) Region-Based Editing (c) Cross-Canvas Mashup

Figure 5: System features. (a) Collage-Based Canvas—compose AI-generated/imported fragments. (b) Region-Based Editing—re-prompt and refine selected regions in place as a new fragment. (c) Cross-Canvas Mashup— recombine ideas across canvases through reuse or image mashup feature.

5.1 Participants and Procedure

We recruited 24 participants (18 female, 6 male) via online postings at our institution, nearby universities, and industry and research institutes. Participants had a mean age of 25.4 years (SD = 3.3) and came from diverse academic backgrounds, primarily design, as well as business, engineering, computing, and interdisciplinary programs. All reported prior design or HCI-related experience. Detailed demographics are shown in the Appendix (§A).

The study procedure is shown in Fig. 6. The study was conducted remotely, with participants joining via video conferencing and using the web-based system on personal devices, reflecting real-world remote collaboration. Participants completed a visual brainwriting task using either AtoMix or a baseline system on the topic *"Design experiences for a learning environment in virtual reality"*. The study sessions lasted 1.5 hours, and participants were compensated approximately $30 USD.

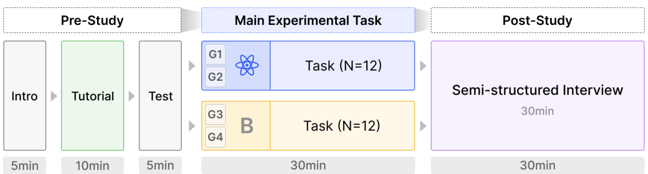

Figure 6: User study protocol. The remote study consisted of a 30-minute ideation phase using either the baseline system or AtoMix, followed by a 30-minute semi-structured group interview.

To isolate AtoMix's design effects—integrated generation, granular editing, and cross-canvas recombination—we developed a baseline system that retained the core visual brainwriting structure while omitting these features. Both systems shared the same multi-canvas setup, user assignment, and 6–3–5 round-based automation, ensuring differences stemmed from interaction design rather than collaboration mechanics. baseline participants used a separate image generator with the same model (Gemini 2.5 Flash Image), generating complete images that were copied onto their canvases, reflecting a conventional GenAI-assisted workflow. In contrast, AtoMix participants completed the same six-round session directly within the deployed web application.

For each group, participants first completed a 10-minute tutorial and a 5-minute practice round to familiarize themselves with system functions (i.e., image generation and modification, collaging, and layer manipulation). They then ideated for 30 minutes on the given topic using the 6–3–5 brainwriting framework, generating three initial ideas and passing their canvas every 5 minutes for six rounds, resulting in a cumulative set of ideas across participants.

5.2 Measures

After the task, participants took part in a 30-minute semi-structured **group interview** to reflect on their experience with the assigned system, including how it shaped their brainwriting flow and how features such as collaging, image modification, and cross-canvas mashup influenced ideation. We also collected **user logs** capturing all canvas interactions (i.e., prompting, image editing, tool use, and canvas mashup) with timestamps, as well as prompt submissions and generated outputs. For the AtoMix condition, we additionally logged region-based edits, fragment reuse, and cross-canvas mashup requests. Screen recordings provided further interaction context. To assess output characteristics, we collected **canvas visual outputs**, consisting of the final artifacts from each session. Each session produced 108 canvases (6 papers × 6 rounds × 3 idea slots), enabling analysis of compositional expressiveness, iterative visual development across rounds, and idea differentiation within rows.

5.3 Analysis

5.3.1 Action Log Analysis. We logged user interactions at the event level and defined five action categories: *Prompting*, *Tool*, *Text Annotation*, *Refine* and *Cross-pollination*. *Prompting* captured text inputs used to generate images (via the integrated generator in AtoMix or an external interface in Baseline). *Tool* included non-generative edits such as background removal, drawing, text annotation, lasso-based selection and movement, and layer reordering. *Text Annotation* recorded text description updates beneath each canvases. Two categories were exclusive to AtoMix. *Refine* captured region-based AI edits using lasso selection and logged an associated prompting event. *Cross-pollination* captured cross-canvas reuse, including copying images from others' canvases and synthesizing new images from multiple sources. For each event, we recorded participant ID, round, and condition, and aggregated counts by participant, round, and system for analysis.

5.3.2 LLM-as-Judge. We used a vision-enabled Large Language Model (LLM) as an automated judge to evaluate visual ideation outputs at scale, enabling consistent, fine-grained assessment [23, 35]. We employed claude-sonnet-4-5 to assess two dimensions of ideation output. First, *iterative visual development* captured how ideas evolved across rounds within each column (6 canvases representing sequential contributions to one idea slot). The LLM rated visual continuity (maintenance of recognizable elements across rounds) and idea progression (development, expansion, or deepening of concepts) on 5-point scales and classified the dominant development pattern (i.e., linear, branching, iterative, discontinuous). Second, *horizontal idea differentiation* captured diversity within each row (3 canvases representing distinct ideas from the same round). The LLM rated visual distinctiveness and conceptual diversity on 5-point scales and flagged redundancy when ideas lacked meaningful differentiation (§B).

For each evaluation, we used structured prompts, anchored rating scales with examples for each score level, and JSON-formatted outputs, and were blind to system condition. A Python script automated image submission, response parsing, and score aggregation. We analyzed 72 vertical development cases (3 columns × 6 papers × 4 groups) and 144 horizontal differentiation cases (6 rows × 6 papers × 4 groups). To validate LLM judgments, one researcher re-coded a random 20% subset using the same rubrics, yielding high agreement (92.7% for visual continuity; 89.2% for conceptual diversity), consistent with standard reliability thresholds. We acknowledge that LLM-based evaluation may carry systematic biases that a single-coder spot check cannot fully detect; we therefore treat these scores as complementary to, rather than a replacement for, our qualitative and log-based analyses.

5.3.3 Thematic Analysis. To examine user perceptions of AtoMix (RQ3), we analyzed interview transcripts using a general inductive approach [52]. Two researchers independently coded the data, focusing on perceived strengths and limitations of AtoMix for brainwriting. One coder proposed 11 initial themes, which were refined by the second coder. Both then reapplied the finalized codes, achieving strong inter-rater reliability (0.80) [31]. Discrepancies were resolved through discussion, and the findings were organized into themes, highlighting key insights into user experiences with AtoMix.

6 Results

6.1 RQ1. How do ideation artifacts differ between AtoMix and the baselinebaseline system?

We examined the visual ideation artifacts produced during the brainwriting activity, focusing on how system affordances shaped their form and evolution. We analyze artifact-level characteristics that capture how participants externalized, developed, and differentiated ideas across rounds. Our analysis integrates qualitative observations, interaction logs, and visual model–based evaluations, revealing four key patterns that distinguish artifacts produced with each system.

6.1.1 Compositional Expressiveness. We quantified compositional expressiveness as the number of images added to each canvas, computed from interaction logs across all rounds. Participants using AtoMix added significantly more images per canvas (M=9.5, SD=3.26) than baseline participants (M=4.42, SD=3.06, $p<.001^{***}$) across the six rounds. This difference was consistent across all six rounds. Figure 8 shows the distribution of image counts per canvas, with higher densities in the AtoMix condition than in the baseline.

Post-study interviews show that participants in the AtoMix condition frequently placed multiple generated images on a single canvas. Six participants described workflows involving generating, layering, and arranging elements to iteratively build a concept: "I generated new images with the prompt and then combined them" (P2). Three participants used a background-first strategy, adding details incrementally (P1). In contrast, baseline participants engaged in limited collaging, often defaulting to a single image because they did not consider combining multiple images (P21, P22). While some participants added multiple images, these were typically placed side by side rather than composed into a unified visual: "I added additional images with similar concepts," but "didn't compose one big image from two" (P20).

6.1.2 Iterative Visual Development. We examined iterative visual development by assessing how ideas evolved across rounds within each column (Figure 9a). LLM evaluations indicated higher visual continuity (M=3.72 vs. 3.44) and idea progression (M=3.56 vs. 3.19) compared to the baseline; however, these differences were modest and did not reach statistical significance (Table 1), so they should be interpreted as directional trends rather than confirmed effects. Examples of iterative visual development patterns are shown in Figure 7 (2).

Post-study interviews indicated that participants using AtoMix actively built on prior ideas by scanning existing work and adding elements intended for later integration. As P2 noted, "I just added similar-feeling elements and thought maybe the next person would combine them." This often produced a relay- or story-like flow, with each participant extending the previous contribution (P7). While some felt pressure to converge (P7), others experienced this shared narrative as collaborative momentum—"doing it together because there was a story" (P10). In contrast, iteration in the baseline condition was less structured and cumulative. Early rounds were highly divergent, with participants generating ideas independently rather than building on prior work (P17). Although some later referenced earlier ideas (P14), many struggled to sustain coherence, with core ideas becoming "diluted over time" (P14, P15, P18).

6.1.3 Horizontal Idea Differentiation. We assessed horizontal idea differentiation by evaluating visual distinctiveness and conceptual diversity across parallel idea slots within each row (Figure 9b). LLM-based evaluations showed higher visual distinctiveness (M = 4.61 vs. 3.67, $p < .001^{***}$) and conceptual diversity (M = 4.35 vs. 3.53, $p < .01^{**}$), leading to more divergent ideas in the AtoMix condition compared to the baseline (Table 1). Standard deviations were lower for AtoMix than the baseline for both metrics (distinctiveness: SD = 0.86 vs. 1.71; diversity: SD = 0.98 vs. 1.64).

Interview data indicate that participants in the AtoMix condition sustained divergence by generating multiple ideas in parallel within each round. Some reported distributing effort across canvases rather than fully developing a single idea (i.e., "I couldn't fully develop all three ideas with good quality," P2). Participants also described comparing ideas across canvases (P7) and using the cross-pollination feature to introduce variation in their ideas, as P9 noted, the "lightbulb (image-mashup) feature randomly selected canvases and gave really quirky new ideas."

6.1.4 Canvas Utilization. To examine canvas utilization, we analyzed interaction logs capturing empty canvases left unmodified at the end of each round. The log data show that AtoMix participants left significantly fewer empty canvases (37 out of 216, 17.1%) than baseline participants (62 out of 216, 28.7%) across the six rounds ($p<.01^{**}$).

Post-study interviews indicate that participants in the AtoMix often completed canvases under time pressure by copying, reusing, and extending existing images rather than leaving canvases blank. As one participant noted, "I just copied existing elements and added things on top" (P1), while others described continuing the existing flow through cross-pollination (P5, P8). In contrast, baseline participants frequently left blank or partially completed canvases, often due to time constraints, image generation delays, or workflow fragmentation. Participants described "moving on without finishing images" (P14, P18) due to difficulty generating visuals that matched their intent (P23) or losing time while switching tools (P18).

6.2 RQ2. How do interaction patterns in AtoMix and the baseline system influence ideation artifacts?

Having characterized the outputs, we next examine interaction practices associated with each system. Using interaction logs and

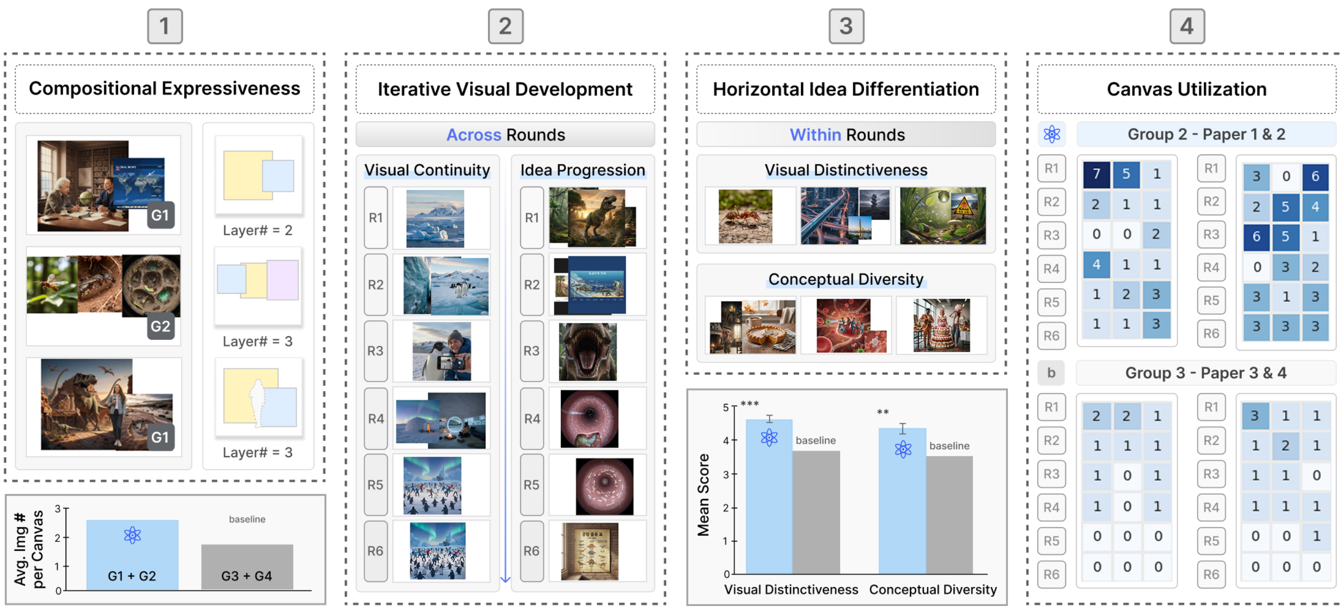

Figure 7: Artifact characteristics of AtoMix. Examples from user study sessions illustrate key output differences. Left: AtoMix produces canvases with higher (1) compositional expressiveness (more images per canvas) and greater (2) iterative visual development (higher visual continuity and idea progression scores). Right: AtoMix yields greater (3) horizontal idea differentiation (higher visual distinctiveness and conceptual diversity scores) and more consistent (4) canvas utilization (fewer empty canvases).

Vertical: Iterative Visual Development				
Metric	AtoMix M (SD)	baseline M (SD)	p-value	Cohen's d
Visual Continuity	3.72 (1.23)	3.44 (1.03)	0.1467 (n.s.)	0.245
Idea Progression	3.56 (1.11)	3.19 (1.17)	0.1678 (n.s.)	0.318
Horizontal: Idea Differentiation				
Visual Distinctiveness	4.61 (0.86)	3.67 (1.71)	0.0007***	0.696
Conceptual Diversity	4.35 (0.98)	3.53 (1.64)	0.0037**	0.608

Table 1: LLM evaluation of vertical (iterative development) and horizontal (idea differentiation) metrics: Mann–Whitney U test results comparing AtoMix and baseline.

post-study interviews, we analyze how participants generated, modified, combined, and communicated ideas across rounds. This analysis identifies three recurring interaction patterns associated with differences in system affordances. Figure 10 shows the full action sequences for each participant across rounds.

6.2.1 Atomic Prompting. We examined prompting behavior using interaction logs capturing prompt counts and lengths across all rounds. The log data show that participants using AtoMix submitted significantly more prompts (M=15.25, SD=5.03) than baseline participants (M=8.58, SD=3.14) across the six rounds ($p<.001^{***}$). Additionally, prompts in the AtoMix condition were shorter on average (M=3.26 words, SD=1.94) compared to baseline (M=10.93 words, SD=7.61) ($p<.001^{***}$). Figure 11 illustrates these differences in prompting behavior.

Post-study interviews data indicate that participants in the AtoMix condition frequently submitted short prompts focused on individual elements as they found detailed prompts ineffective, as P1 noted, "If I input a very long description… it doesn't work or takes too long, so I just input the background only,". This granular approach was especially common early in ideation, when participants "just made things as they came" (P11), enabling rapid exploration. In contrast, baseline participants followed an iterative prompt-refinement pattern, repeatedly editing and reusing a single prompt. Participants described starting with an initial description and incrementally modifying it based on generated outputs (P15); as P16 explained, "I would just take the previous description as my prompt, edit it, generate the image, and bring it over."

6.2.2 Granular Editing. We examined granular editing behavior using interaction logs capturing tool use and image-guided generation.

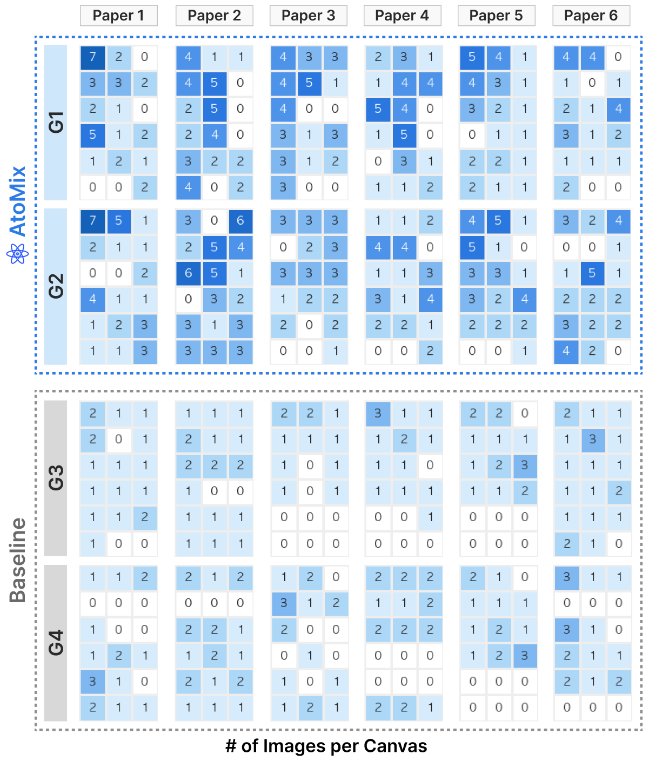

Figure 8: Heatmap of image counts per canvas across all groups. Each square represents a single canvas; there are 18 canvases per paper and 6 papers per brainwriting session, yielding 108 canvases per group in total. Darker colors indicate canvases with more images. AtoMix exhibits consistently higher image counts across participants and rounds compared to the baseline.

AtoMix supports granular editing through tool-based operations (i.e., lasso selection and background removal) and image-guided generation combining text prompts with existing images. Log data show that AtoMix participants used tool-based operations 140 times (104 background removals, 23 lasso selections, and 13 text box uses). In addition, 46 prompts (12% of all AtoMix prompts) used image-guided generation, combining visual references with text instructions.

Post-study interviews indicated that granular editing enabled participants to iteratively build on prior images by targeting specific elements for change. For example, P7 explained, "If I had an idea I could add to the previous one, I would screenshot (lasso) it and say something like 'modify this part.'" In contrast, using the baseline system, participants frequently expressed a desire for granular editing. They noted that being unable to edit existing images or work on the same canvas limited their ability to express and refine ideas: "If I could create on the same page or edit existing images, I think I could express my thoughts better" (P18). Participants reported that this approach affected visual consistency, as "regenerating images from scratch made style and overall image coherence suffer when modifications needed were small" (P23). Reflecting on this gap, they emphasized the need for fine-tuning capabilities once ideas were

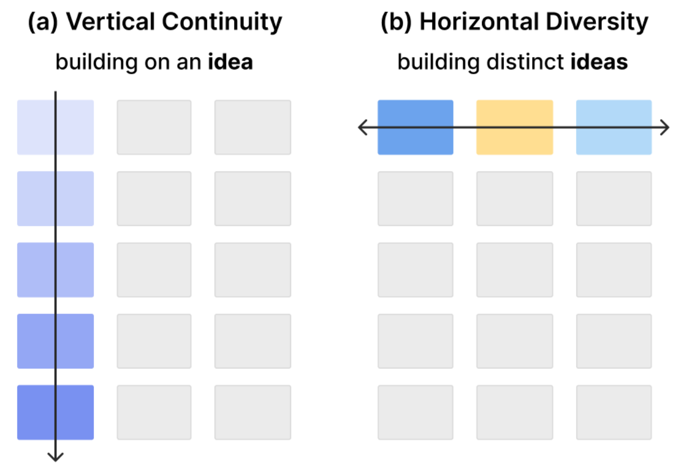

Figure 9: Iterative visual development and horizontal idea differentiation. (a) Vertical continuity captures how ideas evolve across rounds within a single column—measuring visual continuity and idea progression as successive participants build on prior work. (b) Horizontal differentiation captures how ideas vary across parallel canvases within the same round—measuring visual distinctiveness and conceptual diversity among concurrent contributions.

established, such as "fine-tuning options to adjust detailed image elements" or modifying "only specific parts" (P21).

6.2.3 Cross-Pollination Behaviors. We examined cross-pollination behaviors using interaction logs capturing image reuse, references to prior canvases, and cross-canvas operations across rounds. Log data show that participants performed 58 cross-pollination actions in the AtoMix condition, including use of the image mashup feature and object paste actions from other canvases.

Interview data indicate that participants in the AtoMix condition reused and extended content from prior canvases. They treated prior work as references, making incremental additions rather than starting from scratch. Participants described using combination and idea-generation features to mix ideas across boards; for example, "I mixed things from other boards" (P4). Others extended ideas through lightweight reworking, such as copy-pasting an image and adding small elements (P8). In contrast, baseline participants built on others' ideas less directly and with greater effort. They expressed a need for better support to reuse and extend prior work; as P17 noted, "It would be much better if I could utilize previous people's ideas or generated images to supplement my own ideas." Time pressure further limited cross-pollination, with participants feeling compelled to create new content rather than reuse prior work (P15).

6.3 RQ3. How do participants perceive and experience AtoMix for visual brainwriting?

To address RQ3, we analyzed participants' experiences with AtoMix during collaborative visual brainwriting using post-study interviews. Our analysis identified four key themes, which capture how

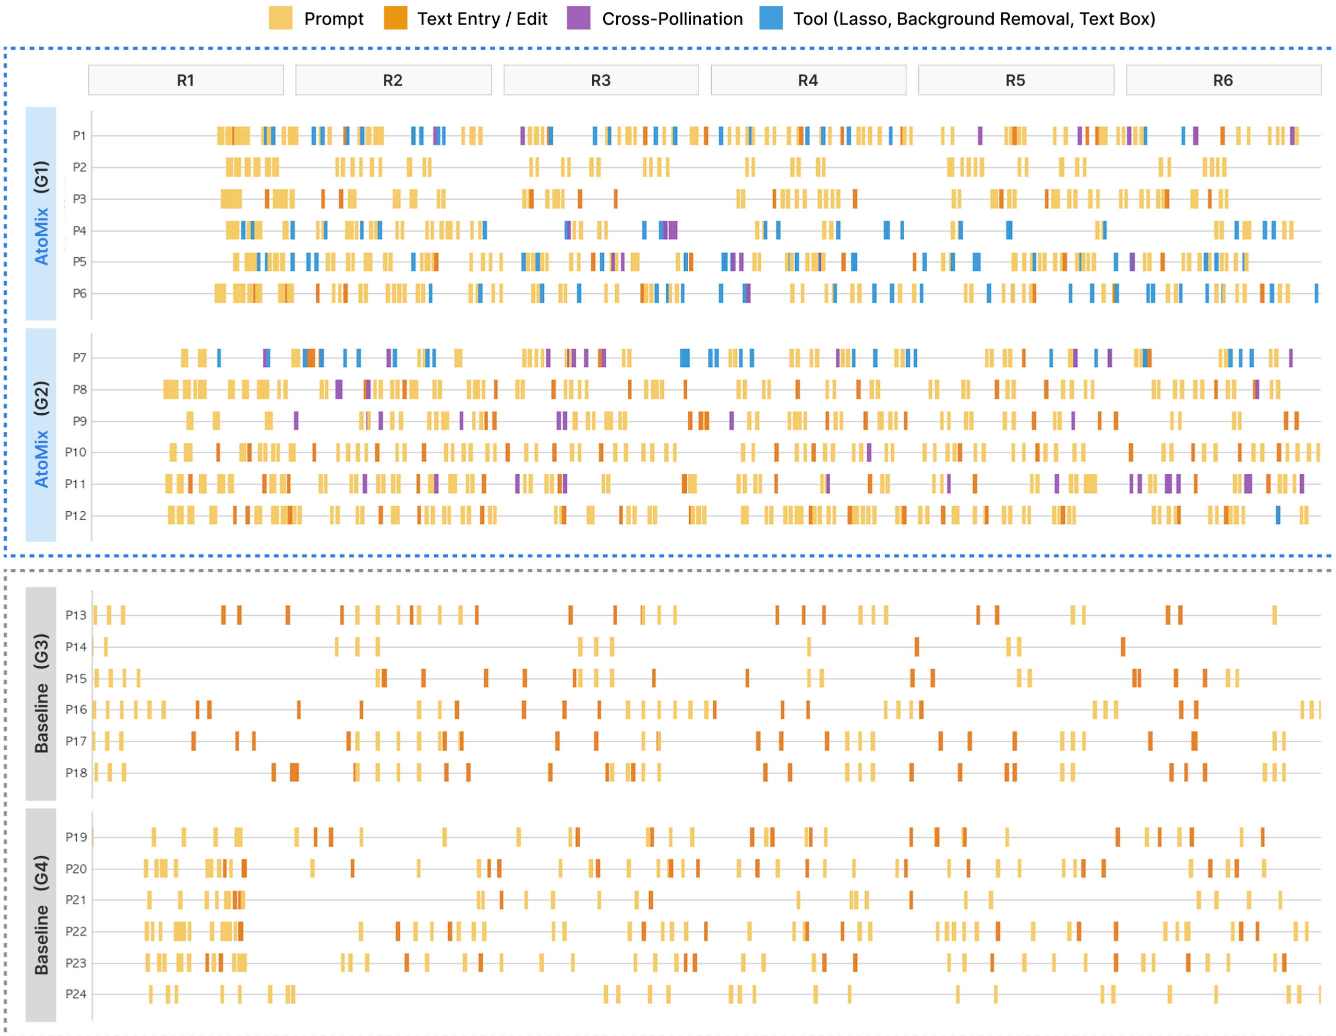

Figure 10: User action sequences across rounds in AtoMix and BASELINE conditions. Each row represents a participant's sequence of actions over six rounds, with colors indicating different action types (i.e., prompt submission, text entry, and tool usage). AtoMix participants exhibit more frequent and varied actions, reflecting higher engagement and interaction complexity compared to BASELINE participants.

AtoMix shaped both collaborative dynamics and participants' subjective experience of ideation.

6.3.1 Workflow Integration. Our thematic analysis revealed that participants valued AtoMix's integrated workflow, which allowed them to search, generate, and combine images within a single environment. Participants appreciated being able to work without switching tools, describing this as "faster and less hectic than relying on external apps" (P2, P7), which supported the rapid generation of multiple ideas (P5). Participants also noted the usefulness of contextual features, such as "viewing the prompts behind generated images to aid interpretation" (P5). Contrarily, in the BASELINE condition, participants described a fragmented and cumbersome workflow that disrupted ideation flow. Moving between writing ideas and generating images was seen as inefficient; "writing ideas then having to move to generate images felt a bit cumbersome" (P18). Participants expressed a desire for tighter integration, suggesting that being able to generate images directly alongside text "on one page" would improve immersion and efficiency (P18). They highlighted friction in reusing prior work, such as having to retype prompts (P21) or being unable to easily build on others' images

6.3.2 Collaborative Experience. Our results revealed that participants using AtoMix frequently described the process as collaborative, cumulative, and story-like (P1, P3, P7, P10, P12). Participants enjoyed seeing ideas build over time, noting that "everything accumulated" and "became more detailed" (P1), and often felt they were responding directly to others' contributions which made collaboration feel like "passing a ball back and forth" (P1). In contrast, BASELINE participants more often experienced the task as individual

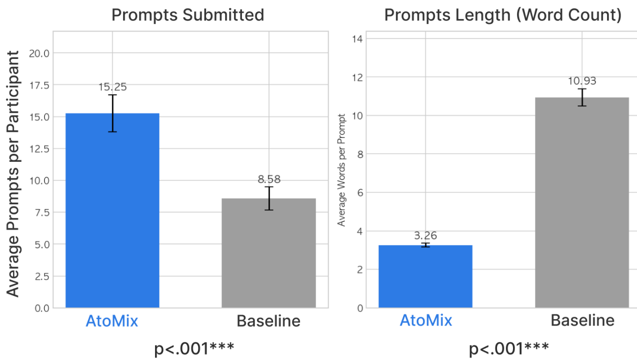

Figure 11: Prompting behavior comparison. (Left) Average Prompts per Participant. (Right) Average Word Count per Prompt. AtoMix users submitted significantly more prompts with shorter average lengths compared to baseline users.

work performed in sequence rather than true co-creation. While some acknowledged moments of collaboration—"especially early divergence" (P19) or when "referencing others' images" (P20) — many reflected that they were "generating ideas individually rather than developing ideas together" (P18), or "continuously performing given tasks" (P16). Differences in personal priorities led to diluted core ideas (P14, P13), and time pressure further reduced awareness of others' presence (P20).

6.3.3 Engagement Through Immersion and Flow. Participants noted that AtoMix generally fostered stronger engagement, with frequently describing the experience as "fun" (P3) and "immersive" (P11, P9). The integrated workflow helped maintain focus, as P3 noted that "search for images and combine ideas within one window supported deeper thinking". Time pressure often enhanced flow, making the process feel "more like a game" and encouraging rapid ideation (P11, P10). In the baseline condition, participants also experienced immersion, particularly due to time constraints, as P20 noted that "having to create three ideas within the 5-minute time constraint really increased immersion". However, engagement was reported to be more fragile and frequently interrupted by workflow friction, such as "going back and forth between image generation and canvas" (P16), retyping prompts, or waiting for images to load (P13). Participants noted that these interruptions made it harder to maintain flow and regain focus once broken (P18).

6.3.4 AI Control and Creative Agency. Our analysis revealed that AtoMix's image quality was sufficient for early ideation but offered limited control. Participants found it effective for rough idea generation (P5, P6) but noted that outputs were at times uniform or "deviated from intent" (P9), making ideas harder to concretize (P10, P12). This lack of control sometimes reduced the creative agency. In the baseline condition, participants also used AI to quickly generate initial visuals, but faced similar control limitations. AI outputs often captured only "about 80% of intent (P14)", lacked text support (P16), and were rarely refined further once acceptable (P19).

Summary of Findings. Across three research questions, our results show that AtoMix produced more compositionally rich canvases with greater horizontal differentiation and fewer empty slots, while vertical continuity showed a positive but non-significant trend. These artifact differences were associated with three interaction patterns: atomic prompting, granular editing, and cross-pollination, enabled by AtoMix's fragment-level affordances. Participants valued the integrated workflow and collaborative momentum, though both conditions faced similar limitations in AI controllability.

7 Discussion

7.1 From Outputs to Materials

Our artifact analysis shows that canvases created in the AtoMix condition differed from those in the baseline along three key dimensions: they featured more compositionally rich collages, showed higher variety of ideas across parallel canvases and increased canvas utilization. We interpret these differences as consequences of reframing generative AI outputs from *finished products* to *manipulable materials*.

The increase in multi-element collages appears linked to AtoMix's affordances—drag-and-drop placement, editing tools, and layering—which reduced the commitment required per image (§6.1.1). In the baseline, each generation was a high-stakes attempt to capture a full concept; partial matches felt like wasted effort. AtoMix reframed partial fragments as useful building blocks, allowing participants to generate a background, then a foreground object, then an annotation—assembling meaning incrementally rather than demanding it from a single prompt. This aligns with Tversky's observation that external representations become thinking tools when manipulated piecemeal [53]. The same low-commitment workflow explains vertical continuity (§6.1.2): when prior contributions can be selectively preserved or recombined rather than regenerated wholesale, participants have structural incentives to build on existing work. Horizontal differentiation (§6.1.3) likewise increased because branching into a new direction cost only a short prompt and a drag operation, leading to more divergent ideas. Even under time pressure, participants could copy, extend, or remix existing fragments rather than leave canvases blank, preserving ideation momentum and reducing the empty canvases that interrupted baseline sessions (§6.1.4).

Across all four dimensions, the common thread is treating generated images as materials that are partial, combinable, and low-cost to modify. This unlocked cumulative creative processes that monolithic generation impedes. That said, AtoMix's integrated workflow also reduces context switching relative to the baseline, so we cannot rule out that some of the observed benefits may stem from integration itself rather than the compositional framing alone.

7.2 Atomicity, Granularity, and Cross-pollination

The interaction patterns help explain *how* the observed artifact differences emerged. Participants learned that short, single-element prompts were more reliable and faster than long, scene-level descriptions; rather than attempting to craft a single optimal prompt, they issued many small requests and assembled the results spatially (§6.2.1). This behavior parallels least-to-most prompting strategies in language models [63] and cognitive chunking in creative problem-solving [9]: decomposing a complex goal into controllable subgoals reduces cognitive load while keeping alternatives open longer. In contrast, the baseline's prompt-refinement loop—edit, regenerate, evaluate, repeat—tended to lock participants into a single conceptual trajectory, whereas atomic prompting supported parallel exploration.

These atomic prompts fed into granular editing, enabling in-place transformation rather than global regeneration (§6.2.2). By lassoing a region and requesting a local change, participants preserved surrounding context and visual coherence, operationalizing Schön's "seeing–moving–seeing" cycle [40]: seeing a partial result, making a situated adjustment, and immediately observing its consequences. In the baseline, even minor revisions required re-describing entire scenes, disrupting continuity and raising the cost of iteration. Over time, cross-pollination increased as a growing critical mass of fragments made recomposition more attractive than generation from scratch (§6.2.3). AtoMix's mashup feature and copy–paste affordances lowered the friction of such recombination, turning the shared canvas into a reservoir of reusable parts.

Together, these interaction patterns suggest a design principle: *reduce the cost of small, reversible moves*. When generation, modification, and recombination are cheap, participants explore more widely, iterate more frequently, and build on each other's work more readily, producing the artifact characteristics observed in RQ1.

7.3 Implications for GenAI-Augmented Ideation

The challenges that GenAI-augmented ideation often faces—intent drift requiring repeated regeneration [8, 48] and early convergence toward generic outputs [15, 43]—may stem less from GenAI itself than from how it is typically integrated. Standard workflows frame each generation as a complete deliverable to accept, reject, or refine—a monolithic interaction that encourages anchoring and amplifies convergence across participants. AtoMix's compositional approach can help mitigate this pattern: by treating generated images as partial materials to be decomposed, combined, and extended, participants introduced variation at the *assembly* stage rather than relying solely on prompt-level divergence. This produced greater vertical continuity alongside horizontal differentiation, suggesting that cumulative build-on and sustained divergence can coexist under low-cost recombination.

This reframing extends prior work on AI-augmented brainwriting, which showed improvements in fluency, elaboration, and divergence–convergence scaffolding in text-based ideation [41]. Our findings indicate that similar benefits can be realized in visual brainwriting, but only when AI outputs are treated not as completed images but as manipulable components. More broadly, these findings suggest that sustaining diversity and supporting user intent in GenAI creativity tools may depend less on model capability alone than on interaction design—specifically, on whether interfaces frame AI outputs as endpoints or as starting points for further transformation.

We note that because both conditions used the 6-3-5 structure, some effects may reflect tool–method interactions—for example, timed rotations may amplify the benefit of low-cost fragment reuse more than they would in a free-form setting. Still, returning to brainwriting's original promise—parallel idea generation, silent exchange, and cumulative build-on [14]—our study suggests that GenAI can amplify rather than undermine these strengths when integrated at the right level of granularity. The 6–3–5 structure's timed rotations and shared sheets were designed to encourage reinterpretation and incremental extension; AtoMix's fragment-based workflow aligns with these principles by making prior contributions visually accessible and easily remixable. Rather than replacing human ideation with GenAI, compositional workflows position AI as a material supplier that accelerates externalization while preserving the collaborative, cumulative dynamics that make visual brainwriting effective. Building on this perspective, future work could personalize compositional workflows by passively identifying individual contributors during sessions [38], enabling systems to adapt recombination suggestions and cross-pollination opportunities to each participant's contribution history and ideation style.

8 Conclusion

We present AtoMix, a GenAI-augmented collaborative visual brainwriting system for remote users that treats generated images as compositional building blocks rather than complete deliverables. This design supports atomic prompting, granular editing, and cross-pollination, addressing bottlenecks identified in our formative study where prompt–regenerate cycles disrupted incremental, collaborative ideation. A comparative study with 24 participants showed that AtoMix enabled the creation of more compositionally expressive canvases with greater idea differentiation, while still supporting smooth, incremental progression. Participants shifted from crafting long prompts to issuing short, iterative atomic prompts that assembled ideas from visual fragments. Region-based editing and cross-canvas mashups further encouraged reuse and co-creation. Together, these findings demonstrate that GenAI can better support creative practice by shifting from monolithic image generation toward compositional assembly—sustaining cumulative idea development in collaborative design ideation.

Acknowledgments

This work was supported by the National Research Foundation (NRF), Korea, under Grant RS-2024-00436931.

References

[1] Suyun Bae, Oh-Hyun Kwon, Senthil Chandrasegaran, and Kwan-Liu Ma. 2020. Spinneret: Aiding Creative Ideation through Non-Obvious Concept Associations. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. doi:10.1145/3313831.3376746 CHI '20.

[2] Margaret A. Boden. 2004. *The Creative Mind: Myths and Mechanisms* (2nd ed.). Routledge, London.

[3] Sebastian Gregor Bouschery, Vera Blazevic, and Frank T. Piller. 2023. AI-Augmented Creativity: Overcoming the Productivity Loss in Brainstorming Groups. *Academy of Management Proceedings* 2023, 1 (Aug. 2023), 11938. doi:10.5465/AMPROC.2023.11938abstract Publisher: Academy of Management.

[4] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. doi:10.1191/1478088706qp063oa

[5] Alice Cai, Steven R Rick, Jennifer L Heyman, Yanxia Zhang, Alexandre Filipowicz, Matthew Hong, Matt Klenk, and Thomas Malone. 2023. DesignAID: Using Generative AI and Semantic Diversity for Design Inspiration. In *Proceedings of The ACM Collective Intelligence Conference* (Delft, Netherlands) *(CI '23)*. Association for Computing Machinery, New York, NY, USA, 1–11. doi:10.1145/3582269.3615596

[6] Jill Cao, Scott D. Fleming, Margaret Burnett, and Christopher Scaffidi. 2015. Idea Garden: Situated Support for Problem Solving by End-User Programmers. *Interacting with Computers* 27, 6 (Nov. 2015), 640–660. doi:10.1093/iwc/iwu022

[7] DaEun Choi, Sumin Hong, Jeongeon Park, John Joon Young Chung, and Juho Kim. 2024. CreativeConnect: Supporting Reference Recombination for Graphic Design Ideation with Generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–25. doi:10.1145/3613904.3642794

[8] Peter Dalsgaard. 2025. Creative Ambiguity and Cognitive Tension in Generative AI Tools. In *Proceedings of the 36th Annual Conference of the European Association of Cognitive Ergonomics (ECCE '25)*. Association for Computing Machinery, New York, NY, USA, Article 4, 10 pages. doi:10.1145/3746175.3746183

[9] Ronald A. Finke, Thomas B. Ward, and Steven M. Smith. 1992. *Creative Cognition: Theory, Research, and Applications*. The MIT Press. doi:10.7551/mitpress/7722.001.0001

[10] Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, and Peter Dalsgaard. 2019. Mapping the Landscape of Creativity Support Tools in HCI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. doi:10.1145/3290605.3300619

[11] Florian Geyer, Jochen Budzinski, and Harald Reiterer. 2012. IdeaVis: a hybrid workspace and interactive visualization for paper-based collaborative sketching sessions. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design*. ACM, Copenhagen Denmark, 331–340. doi:10.1145/2399016.2399069

[12] Joy Paul Guilford. 1950. Creativity. *American psychologist* 5, 9 (1950), 444. doi:10.1037/h0063487

[13] F. G. Haddad, K. M. Jang, F. Duarte, et al. 2025. User-centered evaluation of visual generative AI for city design: An exploratory technology acceptance model analysis. *Cognition, Technology & Work* (2025). doi:10.1007/s10111-025-00816-7

[14] Peter A. Heslin. 2009. Better than brainstorming? Potential contextual boundary conditions to brainwriting for idea generation in organizations. *Journal of Occupational and Organizational Psychology* 82, 1 (2009), 129–145. doi:10.1348/096317908X285642 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1348/096317908X285642.

[15] Niklas Holzner, Sebastian Maier, and Stefan Feuerriegel. 2025. Generative AI and Creativity: A Systematic Literature Review and Meta-Analysis. arXiv:2505.17241 [cs.HC] https://arxiv.org/abs/2505.17241

[16] Chun-Chieh Huang, Tsai-Yen Li, Hao-Chuan Wang, and Chun-Yen Chang. 2007. A Collaborative Support Tool for Creativity Learning: Idea Storming Cube. In *Proceedings of the 7th IEEE International Conference on Advanced Learning Technologies*. doi:10.1109/ICALT.2007.6 ICALT '07.

[17] Youngseung Jeon, Seungwan Jin, Patrick C. Shih, and Kyungsik Han. 2021. FashionQ: An AI-Driven Creativity Support Tool for Facilitating Ideation in Fashion Design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3411764.3445093

[18] Yucheng Jin, Ge Gao, Tiffany Ong, and Xi Su. 2021. CreativeSVG: A Creativity Support Tool for Abstract Background Design with Generative Vector Graphics. In *Joint Proceedings of the ACM IUI 2021 Workshops*. 1–6. https://ceur-ws.org/Vol-2903/IUI21WS-HAIGEN-6.pdf

[19] Hyeonsu B Kang, David Chuan-En Lin, Nikolas Martelaro, Aniket Kittur, Yan-Ying Chen, and Matthew K. Hong. 2024. BioSpark: An End-to-End Generative System for Biological-Analogical Inspirations and Ideation. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3613905.3651035

[20] Rubaiat Habib Kazi, Kien Chuan Chua, Shengdong Zhao, Richard Davis, and Kok-Lim Low. 2011. SandCanvas: a multi-touch art medium inspired by sand animation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. Association for Computing Machinery, New York, NY, USA, 1283–1292. doi:10.1145/1978942.1979133

[21] Hyunyoung Kim, Kuenhwan Kwak, Jongwoo Jung, Insik Myung, and Minsoo Hahn. 2010. Tablaction: collaborative brainstorming system with stylus-fingertip interactions on tablet PCs. In *Proceedings of the 9th ACM SIGGRAPH Conference on Virtual-Reality Continuum and its Applications in Industry*. ACM, Seoul South Korea, 81–84. doi:10.1145/1900179.1900195

[22] Jingoog Kim, Mary Maher, and Safat Siddiqui. 2021. Collaborative Ideation Partner: Design Ideation in Human-AI Co-creativity:. In *Proceedings of the 5th International Conference on Computer-Human Interaction Research and Applications*. SCITEPRESS - Science and Technology Publications, Online Streaming, — Select a Country —, 123–130. doi:10.5220/0010640800003060

[23] Tae Soo Kim, Yoonjoo Lee, Jamin Shin, Young-Ho Kim, and Juho Kim. 2024. EvalLM: Interactive Evaluation of Large Language Model Prompts on User-Defined Criteria. 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 306, 21 pages. doi:10.1145/3613904.3642216

[24] John Knight, Elliot Ross, and Dan Fitton. 2021. Using Ideation Grids to Power Collaborative Creativity in Face-to-Face and Remote Innovation Sessions. In *Creativity - A Force to Innovation*. IntechOpen, 1–18. doi:10.5772/intechopen.93850

[25] Janin Koch, Nicolas Taffin, Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, and Wendy E. Mackay. 2020. ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1 (May 2020), 45:1–45:27. doi:10.1145/3392850

[26] Nicholas Kohn, Paul Paulus, and YunHee Choi. 2011. Building on the ideas of others: An examination of the idea combination process. *Journal of Experimental Social Psychology - J EXP SOC PSYCHOL* 47 (05 2011), 554–561. doi:10.1016/j.jesp.2011.01.004

[27] Jungryun Kwon, Eui-Chul Jung, and Jiheon Kim. 2024. Designer-Generative AI Ideation Process: Generating Images Aligned with Designer Intent in Early-Stage Concept Exploration in Product Design. *Archives of Design Research* 37, 3 (July 2024), 7–23. doi:10.15187/adr.2024.07.37.3.7

[28] Jo'Anne Langham and Neil Paulsen. 2020. Simulating the Cognitive Leap Using Brainwriting. *Design Management Journal* 15, 1 (2020), 85–101. doi:10.1111/dmj.12062 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/dmj.12062.

[29] Hyelip Lee, Seungwoo Je, and Andrea Bianchi. 2017. All4One: A Moderated Sketching Tool for Supporting Idea-Generation with Remote Users. In *Proceedings of the IASDR 2017, 7th International Conference of the International Association of Societies of Design Research*. International Association of Societies of Design Research, Cincinnati, USA. doi:10.7945/C2JH5D

[30] David Chuan-En Lin, Hyeonsu B. Kang, Nikolas Martelaro, Aniket Kittur, Yan-Ying Chen, and Matthew K. Hong. 2025. Inkspire: Supporting Design Exploration with Generative AI through Analogical Sketching. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. doi:10.1145/3706598.3713397

[31] Mary L. McHugh. 2012. Interrater reliability: the kappa statistic. 22, 3 (2012), 276–282.

[32] Amy Melniczuk, Nam Gu Kang, and Shaun Lawson. 2024. Embedding Thinking Strategies within a Tangible Tree to Orchestrate Small Group Brainstorming. In *Proceedings of the Eighteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Cork, Ireland) *(TEI '24)*. Association for Computing Machinery, New York, NY, USA, Article 13, 18 pages. doi:10.1145/3623509.3633361

[33] Amy Melniczuk, Meng Liang, and Julian Preissing. 2025. Exploring Tangible Designs to Improve Interpersonal Connectedness in Remote Group Brainstorming. In *Proceedings of the Nineteenth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '25)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3689050.3704933

[34] Bernard A. Nijstad and Wolfgang Stroebe. 2006. How the Group Affects the Mind: A Cognitive Model of Idea Generation in Groups. *Personality and Social Psychology Review* 10, 3 (2006), 186–213. doi:10.1207/s15327957pspr1003_1 PMID: 16859437.

[35] Qian Pan, Zahra Ashktorab, Michael Desmond, Martin Santillan Cooper, James Johnson, Rahul Nair, Elizabeth Daly, and Werner Geyer. 2024. Human-Centered Design Recommendations for LLM-as-a-Judge. arXiv:2407.03479 [cs.HC] https://arxiv.org/abs/2407.03479

[36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. doi:10.48550/arXiv.2112.10752 arXiv:2112.10752 [cs].

[37] Mark A. Runco and Selcuk Acar. 2012. Divergent Thinking as an Indicator of Creative Potential. *Creativity Research Journal* 24, 1 (2012), 66–75. doi:10.1080/10400419.2012.652929

[38] Alia Saad, Kian Izadi, Anam Ahmad Khan, Pascal Knierim, Stefan Schneegass, Florian Alt, and Yomna Abdelrahman. 2023. HotFoot: Foot-Based User Identification Using Thermal Imaging. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 262, 13 pages. doi:10.1145/3544548.3580924

[39] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. doi:10.48550/arXiv.2205.11487 arXiv:2205.11487 [cs].

[40] Donald A. Schön. 1992. *The Reflective Practitioner: How Professionals Think in Action* (1 ed.). Routledge. doi:10.4324/9781315237473

[41] 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 (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3613904.3642414

[42] Jami J. Shah, Noel Vargas-Hernandez, Joshua D. Summers, and Sanjay Kulkarni. 2001. Collaborative sketching (c-sketch)—An idea generation technique for engineering design. *The Journal of Creative Behavior* 35, 3 (2001), 168–198. doi:10.1002/j.2162-6057.2001.tb01045.x

[43] Joongi Shin, Ankit Khatri, Michael A. Hedderich, Andrés Lucero, and Antti Oulasvirta. 2025. Facilitating Asynchronous Idea Generation and Selection with Chatbots. In *Proceedings of the 36th Australasian Conference on Human-Computer Interaction (OzCHI '24)*. Association for Computing Machinery, New York, NY, USA, 305–323. doi:10.1145/3726986.3726994

[44] Ben Shneiderman. 1983. Direct manipulation: A step beyond programming languages. *Computer* 16, 8 (1983), 57–69. doi:10.1109/MC.1983.1654471
[45] Ben Shneiderman. 2007. Creativity support tools: accelerating discovery and innovation. *Commun. ACM* 50, 12 (Dec. 2007), 20–32. doi:10.1145/1323688.1323689
[46] Pao Siangliulue, Joel Chan, Steven P. Dow, and Krzysztof Z. Gajos. 2016. IdeaHound: Improving Large-scale Collaborative Ideation with Crowd-Powered Real-time Semantic Modeling. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. Association for Computing Machinery, New York, NY, USA, 609–624. doi:10.1145/2984511.2984578
[47] Kihoon Son, DaEun Choi, Tae Soo Kim, Young-Ho Kim, and Juho Kim. 2024. GenQuery: Supporting Expressive Visual Search with Generative Models. 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 180, 19 pages. doi:10.1145/3613904.3642847
[48] Jiaxin Song, Parnian Shahkar, Kate Donahue, and Bhaskar Ray Chaudhury. 2025. Human-AI Collaboration with Misaligned Preferences. In *Proceedings of the 5th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '25)*. Association for Computing Machinery, New York, NY, USA, 294. doi:10.1145/3757887.3767691
[49] Zhida Sun, Zhenyao Zhang, Yue Zhang, Min Lu, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. 2025. Creative Blends of Visual Concepts. doi:10.48550/arXiv.2502.16062 arXiv:2502.16062 [cs].
[50] Robert I. Sutton and Andrew Hargadon. 1996. Brainstorming Groups in Context: Effectiveness in a Product Design Firm. *Administrative Science Quarterly* 41, 4 (1996), 685–718. http://www.jstor.org/stable/2393872
[51] Masaki Suwa and Barbara Tversky. 1997. What do architects and students perceive in their design sketches? A protocol analysis. *Design Studies* 18, 4 (Oct. 1997), 385–403. doi:10.1016/S0142-694X(97)00008-2
[52] David R. Thomas. 2006. A general inductive approach for analyzing qualitative evaluation data. 27, 2 (2006), 237–246. doi:10.1177/1098214005283748 Place: US Publisher: Sage Publications.
[53] Barbara Tversky. 2011. Visualizing Thought. *Topics in Cognitive Science* 3, 3 (2011), 499–535. doi:10.1111/j.1756-8765.2010.01113.x _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1756-8765.2010.01113.x.
[54] Remko Van Der Lugt. 2002. Brainsketching and How it Differs from Brainstorming. *Creativity and Innovation Management* 11, 1 (2002), 43–54. doi:10.1111/1467-8691.00235 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-8691.00235.
[55] Qian Wan and Zhicong Lu. 2023. GANCollage: A GAN-Driven Digital Mood Board to Facilitate Ideation in Creativity Support. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 136–146. doi:10.1145/3563657.3596072
[56] Sitong Wang, Savvas Petridis, Taeahn Kwon, Xiaojuan Ma, and Lydia B Chilton. 2023. PopBlends: Strategies for Conceptual Blending with Large Language Models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3544548.3580948
[57] Wen-Fan Wang, Chien-Ting Lu, Nil Ponsa i Campanyà, Bing-Yu Chen, and Mike Y. Chen. 2025. AIdeation: Designing a Human-AI Collaborative Ideation System for Concept Designers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–28. doi:10.1145/3706598.3714148
[58] Andrew M. Webb. 2015. Pen + Touch Diagramming to Stimulate Design Ideation. In *Proceedings of the ACM SIGCHI Conference on Creativity and Cognition*. 331–332. doi:10.1145/2757226.2764766 C&C '15.
[59] Chauncey Wilson. 2013. Chapter 2 - Brainwriting. In *Brainstorming and Beyond*, Chauncey Wilson (Ed.). Morgan Kaufmann, 43–61. doi:10.1016/B978-0-12-407157-5.00002-6
[60] Chengzhi Zhang, Weijie Wang, Paul Pangaro, Nikolas Martelaro, and Daragh Byrne. 2023. Generative Image AI Using Design Sketches as input: Opportunities and Challenges. In *Proceedings of the 15th Conference on Creativity and Cognition (C&C '23)*. Association for Computing Machinery, New York, NY, USA, 254–261. doi:10.1145/3591196.3596820
[61] 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, 1–22. doi:10.1145/3706598.3713289
[62] Zhenpeng Zhao, Sriram Karthik Badam, Senthil Chandrasegaran, Deok Gun Park, Niklas Elmqvist, Lorraine Kisselburgh, and Karthik Ramani. 2014. skWiki: A Multimedia Sketching System for Collaborative Creativity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1235–1244. doi:10.1145/2556288.2557394
[63] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. 2023. Least-to-Most Prompting Enables Complex Reasoning in Large Language Models. arXiv:2205.10625 [cs.AI] https://arxiv.org/abs/2205.10625

A Participant Demographics

Gender	Male: 6, Female: 18
Age	Mean = 25.4, SD = 3.3
Majors	Industrial Design (10)
	Spatial Design (5)
	Culture Design Management (1)
	Business Administration (1)
	Computer Science and IT Convergence (2)
	Graduate School of Culture Technology (2)
	Mechanical Engineering (1), Bio-Convergence (1),
	School of Transdisciplinary Studies (1)

Table 2: Participant demographics (N=24).

B Task 1: Iterative Visual Development (Vertical Continuity)

For each column of canvases (6 rounds representing sequential contributions to one idea slot), we evaluated visual continuity and idea progression using the following prompt:

Vertical Continuity Evaluation Prompt

```
You are evaluating visual idea development in a collaborative 6-3-5 brainwriting session. You will see 6 images representing the sequentially developing idea slot (vertically) across 6 sequential rounds, where different participants built upon previous contributions.
```

Task: Analyze how the visual idea evolves from Round 1 to Round 6.

Important: Handling Empty Canvases

Some canvases may be EMPTY (blank, no content). Empty canvases indicate that a participant did not contribute to that round. When evaluating, consider how gaps affect the overall narrative and whether later rounds successfully recover or build upon available content.

Evaluation Criteria:

1. Visual Continuity (1-5)

How well do later rounds maintain recognizable visual elements from earlier rounds?

- 1: No connection; each round appears unrelated to others; the visual thread is completely lost
- 2: Weak connection; only superficial or accidental similarities; development chain is broken
- 3: Moderate connection; some elements carried forward but inconsistently; partial continuity
- 4: Strong connection; clear visual thread with intentional variations throughout
- 5: Excellent continuity; cohesive evolution with consistent style/elements; seamless progression

2. Idea Progression (1-5)

Does the concept develop, expand, or deepen across rounds?

- 1: No meaningful progression; ideas appear static, random, or disconnected across rounds
- 2: Minimal progression; only superficial changes; limited evidence of building upon prior work

- 3: Moderate progression; some meaningful development visible, though uneven or interrupted
- 4: Strong progression; clear elaboration and refinement; ideas grow in complexity or detail
- 5: Excellent progression; systematic building with increasing sophistication; rich cumulative development

3. Continuity Analysis
Identify any gaps in the development chain:

- `has_gaps`: true/false
- `gap_locations`: List of round numbers with empty canvases
- `recovery_quality`: If gaps exist—"strong", "partial", "weak", or "n/a"

4. Development Type (categorical)
Classify the dominant pattern: Linear, Branching, Iterative, Discontinuous, Convergent, or Fragmented.

Output Format: JSON with fields for continuity_analysis, visual_continuity (score, reasoning), idea_progression (score, reasoning), development_type, key_observations, strongest_development, and weakest_development.

C Task 2: Horizontal Idea Differentiation (Within-Row Diversity)

For each row of canvases (3 canvases representing distinct ideas from the same round), we evaluated visual distinctiveness and conceptual diversity using the following prompt:

Horizontal Diversity Evaluation Prompt

`You are evaluating idea diversity in a brainstorming session. You will see 3 images created by the same participant in the same round. Each image should represent a DIFFERENT idea for the same design challenge.`

Task: Assess how distinct these 3 ideas are from each other.

Evaluation Criteria:

1. Visual Distinctiveness (1-5)
How visually different are the 3 canvases?

- 1: Nearly identical; same composition/elements with trivial differences
- 2: Low distinctiveness; same visual approach with minor variations
- 3: Moderate distinctiveness; some shared elements but different arrangements
- 4: High distinctiveness; clearly different visual approaches
- 5: Maximum distinctiveness; completely different visual styles/compositions

2. Conceptual Diversity (1-5)
How different are the underlying ideas/concepts?

- 1: Same idea expressed 3 times (redundant)
- 2: Variations of one core idea (e.g., same feature, different settings)
- 3: Related but distinct concepts (e.g., same theme, different features)
- 4: Different concepts with some thematic connection
- 5: Completely different conceptual approaches

3. Redundancy Detection
Identify any redundancy:

- "None": All 3 ideas are meaningfully different
- "Partial": 2 ideas overlap significantly, 1 is distinct
- "High": All 3 ideas are essentially the same concept

Output Format: JSON with fields for visual_distinctiveness (score, reasoning), conceptual_diversity (score, reasoning), redundancy, idea_summaries (3 brief descriptions), most_distinct_pair, and most_similar_pair.