# Co-Ideation Across Time: Revitalizing Legacy Design Sketchnotes with Conversational AI Agents to Foster Intergenerational Collaboration


Yuqing Lucy Li*
MIT Media Lab
Massachusetts Institute of Technology
Cambridge, MA, USA
lucy_li@media.mit.edu

Quincy Kuang*
MIT Media Lab
Massachusetts Institute of Technology
Cambridge, MA, USA
quincyku@media.mit.edu

Xiao Xiao
Institute for Future Technologies
De Vinci Higher Education
Courbevoie, France
MIT Media Lab
Massachusetts Institute of Technology
Cambridge, MA, USA
xiao.xiao@devinci.fr

Jean-Baptiste Labrune
MIT Media Lab
Massachusetts Institute of Technology
Cambridge, MA, USA
labrune@media.mit.edu

Hiroshi Ishii
MIT Media Lab
Massachusetts Institute of Technology
Cambridge, MA, USA
ishii@media.mit.edu


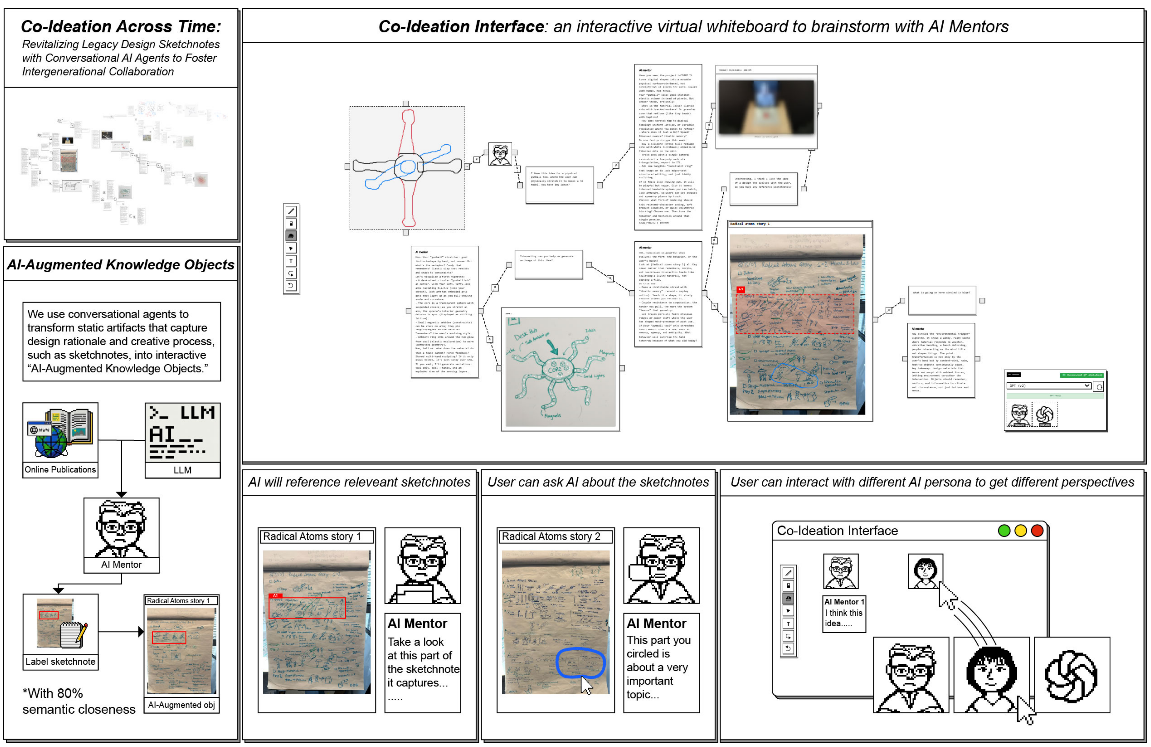


**Figure 1: Overview of the system. The Co-ideation Across Time Interface leverages AI Mentors and AI-augmented Knowledge Objects to support intergenerational collaboration.**

*Both authors contributed equally to this research.



*CHI '26, Barcelona, Spain*



## Abstract

While legacy sketchnotes capture rich design rationales and inspirations, they are rarely reused in contemporary practice. We




ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791258

present Co-Ideation Across Time (CIAT), utilizing Large Language Models (LLMs) to transform decades-old design sketchnotes into interactive "AI-augmented Knowledge Objects". Our system digitizes over 2,000 pages of alumni sketchnotes and connects them with conversational agents trained on corresponding theses and publications, enabling current and future students to engage in multimodal dialogue with past ideas and researchers. An exploratory evaluation with 12 participants showed that interacting with the system stimulated deeper understanding of abstract concepts, idea diversity, and fostered a stronger sense of continuity with the community's legacy. Our contributions are threefold: (1) a method for integrating design legacies with LLM-driven conversational agents; (2) an empirical study demonstrating how this approach supports learning and intergenerational knowledge sharing; and (3) a conceptual framing of AI-Augmented Knowledge Objects as active participants in design ideation.

## CCS Concepts

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

## Keywords

Creativity, Collaboration, Cross-generation, Knowledge Sharing, AI Expert Persona, Sketchnotes, Computer Mediated Communication

**ACM Reference Format:**
Yuqing Lucy Li, Quincy Kuang, Xiao Xiao, Jean-Baptiste Labrune, and Hiroshi Ishii. 2026. Co-Ideation Across Time: Revitalizing Legacy Design Sketchnotes with Conversational AI Agents to Foster Intergenerational Collaboration. 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, 17 pages. https://doi.org/10.1145/3772318.3791258

## 1 Introduction

"If I have seen further, it is by standing on the shoulders of giants."
— Isaac Newton

Sketchnote is a form of visual thinking that combines sketches and written words. Creatives often use sketchnotes as the first instance to prototype ideas. Depending on the project, they capture a wide range of multifaceted considerations, such as functionality, aesthetics, practicality, motivation, user capabilities, or sociopolitical concerns [1, 20]. However, sketchnotes can be difficult to read, handwriting is challenging to decipher, and there is no clear direction of reading order. Therefore, to support knowledge exchange via legacy sketchnotes, we created a conversational agent, called AI Mentor, that is created based on human mentor's publications, theses, talks, and social media content. The conversational agent can answer questions and transform sketchnotes into interactive AI-augmented Knowledge Objects that illuminate past ideation processes for current and future students.

Our paper first presents a pipeline for creating an AI Mentor that can enhance comprehension of messy design legacies. Secondly, we demonstrate the Co-Ideation Across Time (CIAT) interface, which integrates the AI Mentor and sketchnotes to deepen the understanding of established concepts and facilitates brainstorming through intergenerational knowledge transfer. 3) We evaluate the system with an empirical study on what strategies students use when learning and ideating with AI-augmented Knowledge Objects, specifically legacy sketchnotes. Our study focuses on the following research questions:

- **RQ1:** How does a custom conversational agent (AI Mentor) affect comprehension of design legacies and knowledge transfer of ideation processes? What interaction behaviors emerge when participants interact with or without the AI Mentor during the interpretation of sketchnotes?
- **RQ2:** How does CIAT system support users during ideation and brainstorming? How do participants interpret, adapt, or reject the AI Mentor's input during co-ideation?
- **RQ3:** What experiences and social behavior emerge when individuals learn and ideate with CIAT? How does a participant's existing relationship to the human mentor affect the experience with the AI Mentor?

The exploratory evaluation (n = 12) consists of four tasks. Each is designed to give insights into interpreting sketchnotes with and without AI Mentor, how learners interact with an expert-perspective LLM during ideation, and the social or emotional experience of co-ideating across time. The data was collected via transcripts, self-report surveys, interaction logs, and semi-structured interviews. Our main contribution presents empirical data on the experience of interacting with past design sketchnotes using a conversational agent trained on specific knowledge. While we found that sketchnote by itself also helped participants deepen concepts, many were unable to get further confirmation on their interpretations. Whereas, with an AI Mentor that helped enhance readability and reconstruct spatial-temporal relations within a sketchnote, subjects became more inquisitive and allowed them to better understand the rationale behind the drawings and annotations. Regarding ideation, the clear differences in response between an AI Mentor and a general LLM provided participants with a nuanced lens to adapt the generated feedback. This allowed them to explore completely novel directions while maintaining control over the outcome. We also uncovered unexpected insights into how individuals perceived the experience beyond our initial research questions. Many participants reflected that they would be hesitant to ask a human collaborator the same foundational questions, likely due to social pressures. Some even felt encouraged to see that the ideation process from researchers they admire are not much different from their own. Our study contributes to the development of AI-supported intergenerational knowledge sharing.

## 2 Related Work

### 2.1 Ideating with Design Examples

Cognitive research indicates that people enhance their creative ideas by learning from examples, recombining concepts, and iterating to improve them [14, 21, 37, 52]. This foundational principle is the cornerstone of many widely used commercial tools (e.g., Pinterest, Miro), which are designed as collaborative platforms for collecting and sharing design examples, thereby facilitating this cycle of example, iteration, and recombination. Building on this success, HCI researchers are further improving how to curate and structure information to support creative work. Dynamic information retrieval (IR) systems study how to help people find information more effectively [18]. Example-based ideation systems study when (e.g., before, during, or after users form initial ideas) to show which (e.g., visually, semantically, conceptually relevant or distant) examples [20]. While most studies demonstrate the effectiveness of ideation with the right design examples [46], they also agree on one central challenge: Design Fixation [25]. Researchers report close examples tend to result in people copying surface features and limiting their thinking to ideas around shown solutions. Whereas, distant or cross-domain examples can be novel but also misleading, resulting in impractical ideas with no deeper exploration. We observe that a potential problem with exposing people to examples is that they tend to anchor people's attention to finalized projects, while the designers are still in the process. By showing examples of process (sketches, iterations, workflows) instead, we invite creatives to think more about methods, strategies, and transformations.

## 2.2 Capturing Design Process

gIBIS, developed in 1988, uses hypertext to point to design rationale [10]. By design rationale, the authors mean the design problems, alternative resolutions (including those that are later rejected), tradeoff analysis among these alternatives, and a record of tentative and firm commitments made as the problem was discussed and resolved. gIBIS suggests that the computer is a powerful medium for collaboration in understanding the internal structure of design decisions. More recent examples include Sterman et al.'s research on calling for design tools to focus on documentation rather than only creation for HCI education [47], stating similar to how writers store drafts or how ceramicists make work visible in a space, the value of in-progress work can support and encourage norms that value process rather than only outcomes. The annotated portfolios method by Gaver and Bowers argue that neither dimensions of concern nor designers' orientations can be read directly from finished artifacts themselves [16]. Instead, they require annotation to provide a means of capturing a design history for a future audience, uncovering underlying values, and communicating insights and learnings to a wider audience. Related work on digital heritage demonstrates how interactive technologies facilitate collaborative interpretation and meaning-making of cultural artifacts [9]. Recent work further demonstrates how people co-create cultural-heritage stories with generative AI, highlighting both the creative potential and the negotiation of authenticity in AI-mediated cultural narration [19]. Sketchnotes have been studied in HCI as an effective medium for notetaking and enhancing the cognitive process [55], both for creative activities like visual ideation and cognitive tasks such as understanding engineering concepts. They are often drawn and written in real time, capturing the essence of a lecture, a talk, or the serendipity of an idea [48]. Unlike common notes, sketchnotes often emphasize content and ideas, and do not strive for the quality of the final representation [41]. Some researchers believe that these unique attributes of sketchnotes, compared to typeset final pieces, capture tacit knowledge. Macdonald calls it the "scars of the thinking process" [36] and Ishii "Traces of Physical Presence" [22]. However, the current use of sketchnotes beyond serving as records is limited since they can be difficult to interpret [41]. They are static and simplified in nature, often omitting context or relying on personal visual shorthand that others may not easily understand [15, 42].

## 2.3 Ideating with AI Agents

In recent years, the role of generative AI technologies in creative processes has been actively studied [1, 6, 43, 51], demonstrating that, as collaborative tools, they consistently support divergent and convergent thinking, improve idea generation, and framing, especially for novices [49]. Particularly interesting for ideation is how LLMs excel at cross-referencing and generating associative knowledge from a sea of data that would otherwise be inaccessible [2]. Elicit, for example, is a commercial AI service for systematic reviews and meta-analyses of research, with over 400,000 monthly active users. Furthermore, when provided with sufficient context, LLMs can plausibly reconstruct and generate missing parts, such as the Content-Aware Fill function in Adobe Photoshop. Beyond filling missing pixels, the Reframer [33] studied how AI can expand on incomplete sketches, half-formed thoughts, or rough prompts during ideation. Building on this work, we created a custom conversational agent with corresponding knowledge to fill the gaps in understanding sketchnotes. This turns a static sketchnote into an interactive one, allowing users to understand the context of legacy drawings and annotations that would otherwise be extremely challenging to comprehend or be completely lost once the sketchnote author has passed away. In addition to new technical possibilities, Glinka and Müller-Birn show that computational tools can and must scaffold critical-reflective engagement with complex artifacts [17]. Recent ideation tools such as PersonaFlow further demonstrate how LLM-simulated expert perspectives can enhance early-stage conceptual exploration[34]. Research on AI-mediated mentorship similarly emphasizes how AI can provide guidance and skill development [39]. In group settings, Lyu et al. show that AI can act as a sustained teammate, influencing learners' workflows and conceptual strategies[35]. Yet, work in team training highlights that meaningful mentorship requires the coordination and calibration of trust [30].

These are also the weaknesses and risks of AI-mediated mentorship, including over-trust, reduced autonomy, and privacy concerns. Shin et al. found no evidence that LLMs alone improve problem reframing and observed a reduction in agency among less-experienced users [44]. These cautionary findings underscore the need to design AI systems that explicitly support user agency, transparency, and critical engagement rather than passive reliance.

Building on these foundations, our work investigates how learners with foundational familiarity engage with AI-augmented work-in-progress materials, and how such materials shape interpretation, conceptual understanding and ideation.

# 3 AI Pipeline for Intergenerational Knowledge Exchange

To transform legacy design artefacts into accessible knowledge, we designed a tool called Annotation Labeling Interface (see figure 2), consisting of three functions. It is used to create 1) AI Mentors: conversational agents that impersonate human mentors based on related publications, theses, and other materials. 2) AI-augmented

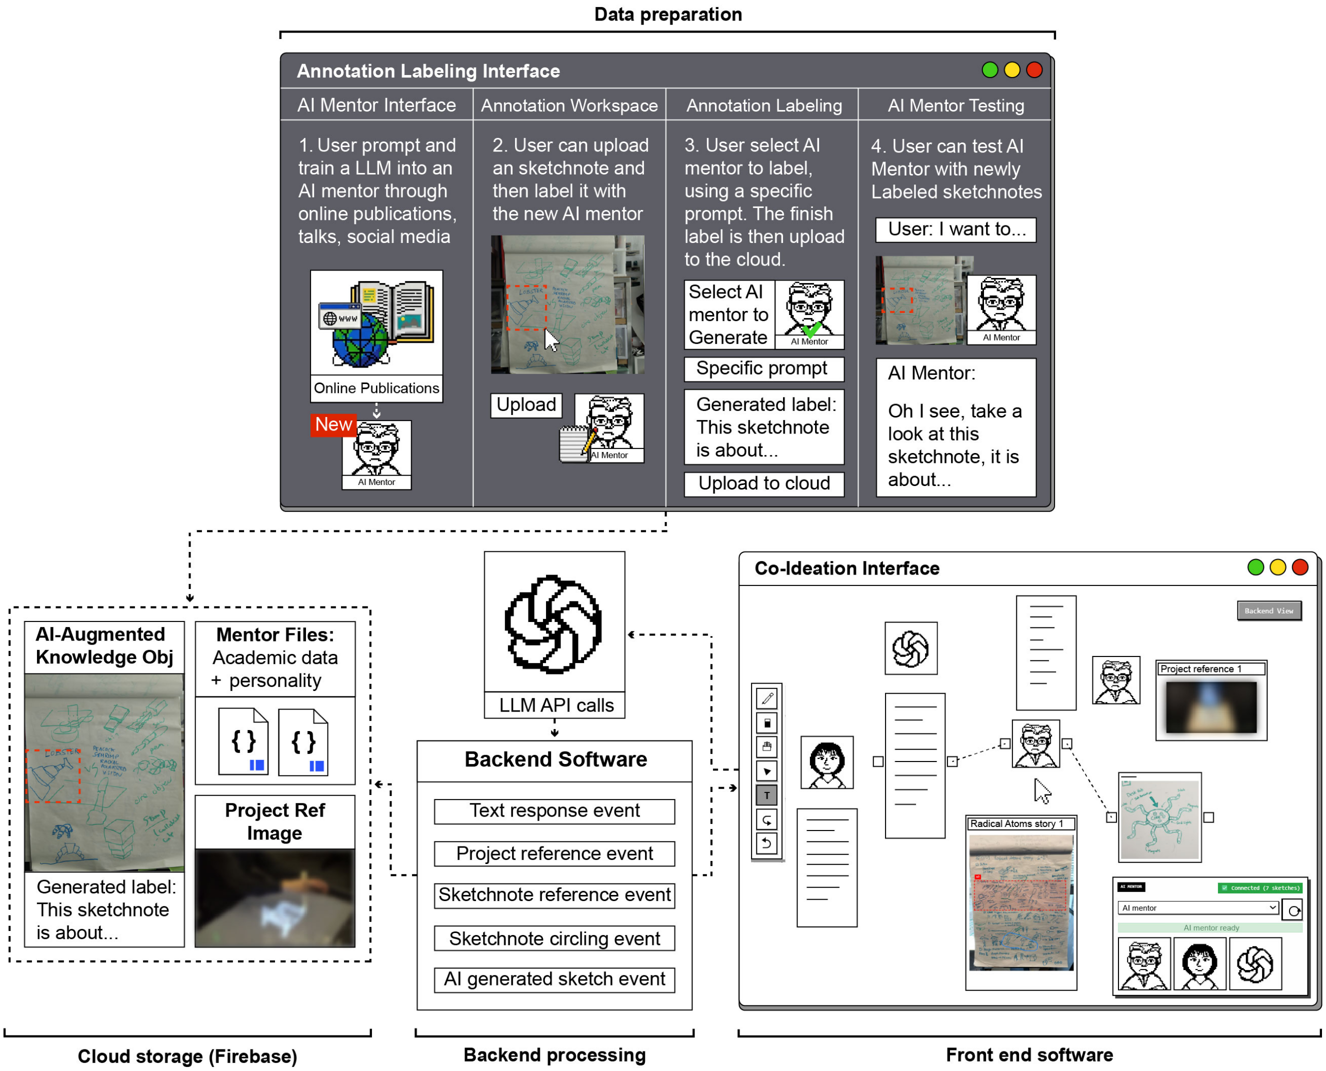


**Figure 2: System pipeline. Users create custom AI Mentors and AI-augmented Knowledge Objects using the Annotation Labeling Interface. The data is then stored on a cloud server and can be retrieved later via API calls.**

Knowledge Objects: legacy design artefacts made useful and effective for brainstorming, 3) to test the quality and accuracy of both AI Mentors and AI-augmented Knowledge Objects

## 3.1 AI Mentors

The input data for the AI Mentor consists of three components: academic publication track-record including images of projects for detailed academic understanding, transcriptions of talks and presentations for high-level comprehension of mentor's motivation and lastly, social media content for the tone of voice. These three components are fed into a RAG (Retrieval-Augmented Generation) system, refined through few-shot prompting and compiled into two editable JSON files called personality and academic data. We did not custom train our model for two reasons. First, the system remains lightweight, easily adjustable, and implementable using natural language [3]. Second, the simplicity and universality of natural language lower the barrier for others to create mentor personas without requiring advanced programming or AI expertise [4]. The finalized AI Mentor with the academic understanding, high-level storytelling and adjusted tone of voice is then stored in cloud storage and can be accessed by any front-end through API calls, see Section 4.

## 3.2 AI-augmented Knowledge Objects

Using the same Annotation Labeling Interface users can then upload legacy design artefacts. While a range of materials could be accommodated with this pipeline, we focused on sketchnotes. With a selection tool users can draw boxes enclosing different regions of the sketchnote. With each selection, an indexed annotation window pops up. One can either manually annotate the region or ask the before created AI Mentor to annotate by posing questions e.g., "How does this drawing in the sketch correlate with xyz concept?", "What is especially interesting about this idea compared to the others?", or "Why did final project not include this aspect?". Users can refine responses or try out different prompts before saving. The system packages the annotation coordinates, text, and metadata into a JSON file alongside the sketchnote image in Firebase. This process converts legacy sketchnotes into an AI-augmented Knowledge Object that is ready for interaction.

After making the AI Mentor and AI-augmented Knowledge Object, users have the option to test the accuracy of both within the same interface.

## 3.3 Accuracy of AI Mentor Responses

To assess the quality of the AI Mentor's response, we conducted a review session with the human mentor, asking them to explain their own sketchnotes. We transcribed their explanations (reference sentences) and compared them to the explanations generated by the AI Mentor created through the data pipeline (candidate sentences). Through evaluating the semantic closeness of our candidate texts against the reference notes using traditional N-gram metrics BLEU [40] and ROUGE [32], as well as BERTScore [11, 54], an NLP metric for contextual embedding similarity. Both BLEU/ROUGE values were low, around F1 = 0.2, indicating a low word overlap and reflecting substantial rewording. However, this is anticipated since the AI Mentor was trained on written language derived from theses, publications, and rehearsed talks, whereas during the explanation sessions, the author, a non-native English speaker, explained more freely. In support, the BERTScore F1 score of 0.8 reveals strong correspondence at the meaning level across all sections. In concurrence with human check, this suggest that, despite stylistic differences, the candidate text can preserve the core meaning and majority of the conceptual content of the reference. Additionally, we observed that our pipeline was able to explain elements on the sketchnote that had never been mentioned in publications. For example, one sketchnote contained a vague scribble with the annotation "Damascus steel knife," which initially appeared irrelevant in the context of materials with physical memory (see figure 3). Yet, upon asking, the AI Mentor gave a general explanation of what Damascus steel is and how in this context the layered steel might have been used as a metaphor to describe how material can have embedded intention, memory, and transformation, like the blade's layered structure. This explanation was approved by the human mentor, who expressed surprise about the LLM's ability. This example illustrates how AI Mentor can uncover latent associations, interpret metaphors, and provide accessible explanations that approximate the original creator's intent. Without a generative AI system, it might be difficult for the AI Mentor to make this connection and answer such specific questions.

This interaction, grounded in the authentic sketchnote, creates a novel brainstorming environment where users' creativity is challenged, inspired, and remixed with ideas and frameworks from the past. Generative AI is essential to our system because it enables open-ended, personalized dialogue and generalization beyond the fixed content of archival sketchnotes. In contrast, retrieval-based approaches are constrained to pre-existing responses and cannot meaningfully address novel or unanticipated user questions.

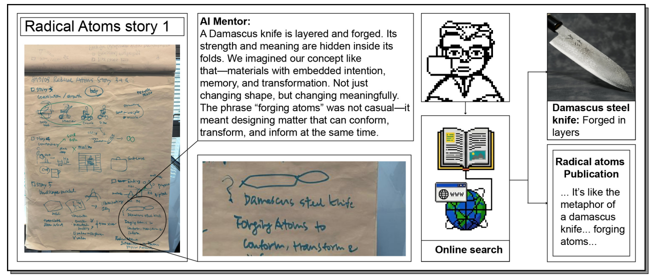


**Figure 3: Damascus steel knife example. By searching for "Damascus steel knife," the AI locates the relevant publication and presents the concept in its correct context, making an otherwise hard-to-understand topic accessible.**

# 4 CIAT Interface

## 4.1 Features

The CIAT Interface is a node-based whiteboard interface that enables users to co-create create with the AI Mentor as well as the AI-augmented Knowledge Object (Figure 4). The goal of the system is to simulate a real-time brainstorming experience with a mentor. First, users can more effectively ask questions about sketchnotes and learn from them because the AI Mentor provides immediate context drawn from the author's related work. Second, instead of receiving answers in a standard sequential text format, users can place LLM responses, media, and generated drawings directly onto the canvas. This spatial layout allows them to engage with the AI Mentor in a way that resembles collaborating with someone on a shared virtual whiteboard. Third, CIAT is designed to let users interact with different AI Mentors by dragging their icons onto the canvas. This will enable users to consult multiple LLM-based expert personas whose expertise aligns with different aspects of one idea, encouraging participants to consider their project from multiple perspectives. Features include:

- **Canvas** A central whiteboard space where users can pan, zoom, and move elements freely.
- **Toolbar menu** Provides tools such as pen and eraser, which allow users to draw or erase both freehand lines and annotations on sketchnotes.
- **Drawing Boxes** Each pen stroke generates a bounding box, and nearby strokes are merged into the same box. Proximal drawings also combine into a single image-based box, which is then sent to the AI for analysis.
- **Connector Dots and Lines** Each element has four connector points that allow users to link nodes. Connected elements are designated as "focused topics," while unconnected ones are still provided as contextual input but not emphasized.
- **AI Mentor Icon** These icons represent the LLM-based expert persona the user is addressing. The user can drag a mentor icon onto the canvas to prompt that specific mentor. By linking a drawing/text box to a mentor icon, the interaction simulates directing a question to an individual mentor.
- **User Text Input Box** Users can drag to create text boxes on the canvas; the text box serves as the main channel for communication with the AI Mentor.

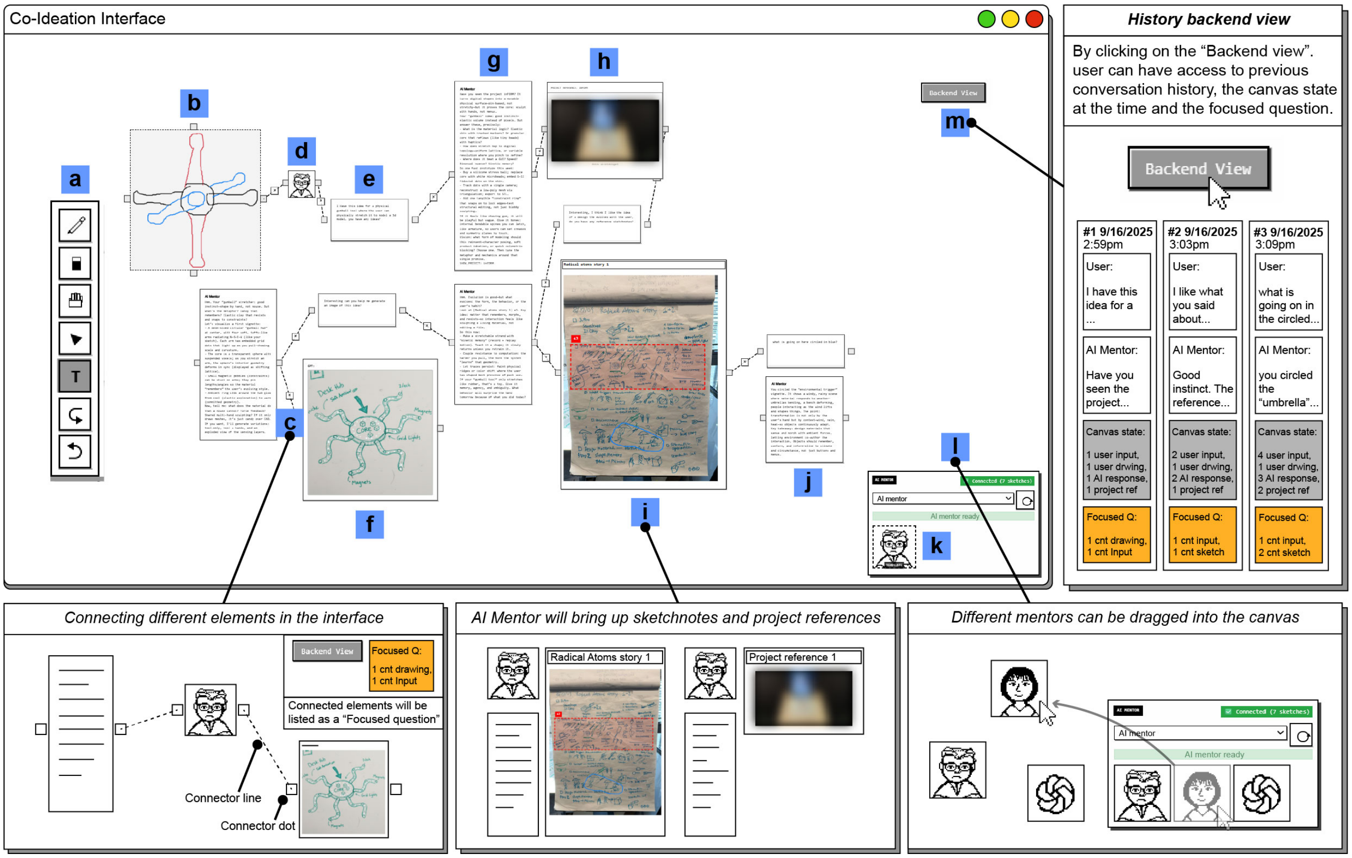


**Figure 4: Co-ideation Interface features break down: (a) tool bar menu, including pen, eraser, move, text box, undo, (b) drawing box, (c) connector dots and lines, (d) AI Mentor icon, (e) user text input box, (f) AI Mentor generated sketch, (g) AI Mentor response box, (h) AI Mentor reference project, (i) sketchnote reference, (j) AI-circled reference response, (k) AI Mentor Icon Template, (l) AI Mentor menu, (m) history view.**

- **AI Mentor Generated Sketch** Visual responses in the style of the mentor's work, used to inspire users by blending historical knowledge with current ideas.
- **AI Mentor Response Box** Contain textual replies generated by the AI Mentor.
- **AI mentor Reference Project** Appear when the LLM identifies a connection with a related past project. This mirrors the way human mentors naturally reference relevant examples in brainstorming.
- **Sketchnote Reference** Generated when the AI Mentor finds a match between a concept brought up by the user with a labeled section of a sketchnote. These include a highlighted excerpt and an explanatory response.
- **AI-Circled Reference Responses** This is the response generated by the AI mentor after a User circle areas of a sketchnote for explanations.
- **AI Mentor Icon Template** A template section that allows users to drag out multiple copies of the respective AI mentor icon.
- **AI Mentor Menu** Users select from a list of available mentors. Once selected, a mentor icon template appears in the menu.
- **History View** A record of past interactions, including full canvas states, enables both users and the AI Mentor to revisit the evolving context of the conversation.

## 4.2 Example Ideation Workflow

A user may begin by sketching or writing their idea on the canvas. The AI Mentor responds by referencing a past project. The user can then ask follow-up questions, prompting the AI Mentor to display process sketchnotes. By circling specific areas, the user requests more detailed explanations, to which the conversational agent responds with context. The user may then ask how these ideas apply to their current project, leading to generated feedback, suggestions, or "homework." The cycle continues, with the AI Mentor offering feedback and generating sketches, until the user feels the brainstorming session has reached a satisfactory conclusion.

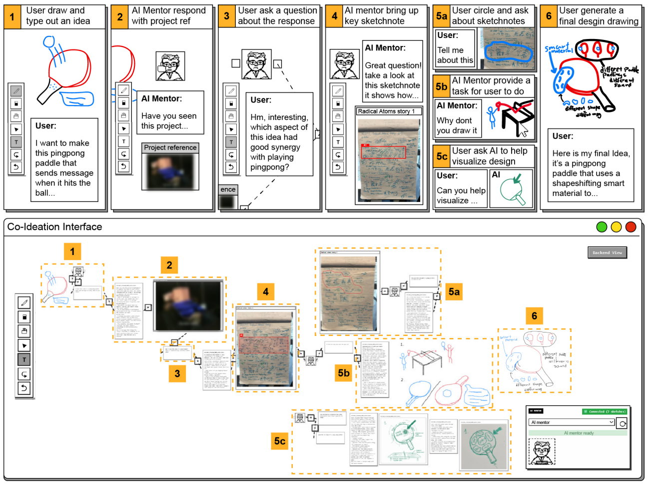


**Figure 5: An example workflow on the Co-Ideation Interface. (1) User may start with a simple idea description and drawing to prompt the AI Mentor, (2) the AI Mentor might then bring up a relevant project for reference (similar to a real conversation with a mentor), (3) user might ask more question about the project, (4) AI Mentor may bring up a relevant process sketchnote related to the project to explain in greater depth, (5a) the user might circle the sketchnote and ask questions, or (5b) the ai mentor might give more task for the user to develop their idea, or (5c) the user might ask AI Mentor to visualize the design. (6) Lastly, the user might draw a final design to conclude the session.**

## 5 Exploratory Evaluation

We design a within-subject exploratory evaluation to examine (1) how interacting with and without the AI Mentor affects the way participants read legacy sketchnotes, (2) observe what strategies they apply when ideating with the AI Mentor in the CIAT interface. In this evaluation, our focus is on the behavior patterns introduced by the AI Mentor, AI-augmented Knowledge Object and CIAT interface, rather than assessing the quality of the ideas generated through it.

### 5.1 Sketchnote Selection

To reduce carry-over effect during the study, we selected 2-sets of sketchnotes based on three criteria: 1. They reflect impactful HCI concepts that are valuable for students to learn. 2. There are at least 10 sketchnotes by the same author on the same topic. 3. The topics have sufficient associated publications and data for creating an AI Mentor.

We contacted researchers across eight topics whose work met these criteria and requested their ideation sketchnotes. From these, we selected "Telepresence"[38] and "Radical Atoms"[23] as the most suitable topics for our study design. The study was approved by the authors' university institutional review board. The following section provides a detailed description of the study procedure.

### 5.2 Procedure

The study took place in a focused room and lasted 90 - 120 minutes. After providing consent and demographic data, participants followed a four-task protocol with fixed durations, conducted in a one-on-one session with a research facilitator who guided the procedure and provided technical support. Participants were encouraged to speak aloud throughout the experience. The entire process was also recorded and shared with 2-3 researchers who took observational notes.

**Task 1: Interpret Telepresence Sketchnote without AI Mentor** Participants first described their baseline understanding of telepresence (pre), then watched a short video of a telepresence-related project, and lastly reviewed that project's work-in-progress sketchnote. After interpreting the sketchnote, participants again explained telepresence (post) and reflected on their experience.

**Task 2: Interpret Radical Atoms Sketchnote with AI Mentor** Participants first described their pre-definition of Radical Atoms (RA), then watched a short Radical Atoms project video, and lastly reviewed the work-in-progress sketchnote of RA. However, this time participants used the CIAT Interface with all its features to review the sketchnote including asking questions to the AI Mentor. After interpreting, participants were again asked to give their post definition of RA.

**Task 3: Create a Radical Atoms-themed Idea with CIAT Interface** Participants were tasked to come up with their own RA–themed idea in the form of a sketch and a brief description. They used the CIAT interface and co-created with the AI Mentor. After completion of the task, participants completed a self-reported survey (Likert items on helpfulness, idea progression, novelty, feeling capable, motivated, creatively challenged, etc.) and a semi-structured interview, see table 2.

**Task 4: Create a Telepresence-themed Idea without CIAT Interface** Participants were asked to brainstorm one Telepresence-themed idea using a commercial whiteboard software. The task followed the same format as Task 3 with an idea sketch and description as outcome. However, the goal is not to compare idea quality between ideating with and without CIAT interface. Instead, the purpose of Task 4 is to understand whether participants were able to carry over learnings and ideation strategies from Task 3 without using the CIAT interface. Therefore, the survey, and brief interview centered around transfer of knowledge. (e.g., "Were there any takeaways from the Co-ideation Interface ideation session that influenced this solo ideation?").

We kept the task order consistent for two reasons. First, the focus of the exploratory evaluation is to capture interaction behavior and patterns that emerge when using this novel interface. To fully understand the difference, we first had participants look at a sketchnote without AI Mentor, as people normally would, to capture their regular interaction behavior. Then we introduced the AI Mentor and asked them to use the tool to understand a sketchnote. Capturing both behaviors in that order allows for comparison of the differences in interaction. Presenting the sketchnote with AI Mentor first might alter how participants normally interact with sketchnotes,

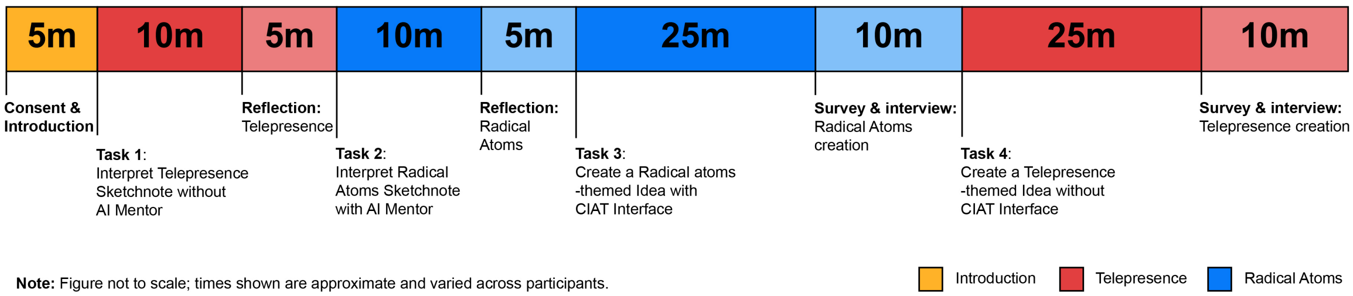


**Figure 6: Study procedure**

| Self-Report Survey |
|---|
| "I feel capable and comfortable of coming up with more Radical Atoms ideas compared to before." |
| "I feel motivated and interested to come up with more Radical Atoms ideas compared to before." |
| "I felt creatively challenged during the task." |
| "I felt my ideas evolved throughout the process." |
| "The sketchnotes felt helpful for someone like me." |
| "I received suggestions that I was not thinking about otherwise." |
| "I felt like I was co-ideating with the interface" |
| "My co-ideator was AI Mentor / Author of Sketchnotes" |

**Table 1: Details of Self-Report Survey**

| Core Interview Questions |
|---|
| Did you go into the task with a clear idea in mind? Did that idea change overtime? Why? |
| How does this ideation process compare to your usual ideation process? |
| Did you feel like you were intellectually engaged / inspired? Can you name an example? |
| What do you or don't you like about looking at other people's sketchnotes? |
| How do you feel about seeing how another person has approached a topic? |
| Who do you feel was in control? Why? |
| How did you perceive the role of the AI? |

**Table 2: Outline of core questions that guided the semi-structured interviews.**

making the comparison more challenging. Second, we acknowledge that the differences in the sketchnote topics themselves may influence the interaction, e.g., participants are more familiar with Telepresence than Radical Atoms. We attempted to mitigate this effect by showing the same project videos before the sketchnotes to ensure a more uniform understanding of the topic among participants. Moreover, we collect their pre- and post-definitions, allowing us to assess the change in understanding rather than the absolute understanding.

## 5.3 Participants

We recruited 12 adult participants (8 male, 4 female), aged 18-34, all with at least four years of design experience. Six of the participants identified as intermediate designers, reporting regular practice and several completed projects, while the other six described themselves as advanced designers with extensive professional or academic experience. All participants reported engaging in ideation more than once per week, with half ideating more than four times per week. Most participants have heard of both Telepresence and Radical Atoms and were able to provide a rough description of the concepts, but they are not confident in stating that they are familiar with them. We deliberately recruited participants with prior understanding of both concepts, as work-in-progress materials can be most useful to individuals with prior familiarity [13, 50], allowing them to comprehend and brainstorm ideas in greater depth.

## 5.4 Data Collection and Analysis

We employed a mixed-methods approach [7] for data collection and analysis. Qualitative data were collected through observations, participants' thinking out loud during tasks, and a semi-structured interview; the latter two were audio-recorded and transcribed for thematic analysis. Quantitative data include users' pre- and post-definitions, interaction logs (e.g., number of questions asked), collected through CIAT backend logs, and self-report Likert-scale surveys. To address RQ1: How people comprehend sketchnotes with and without AI Mentor, we performed a content analysis [29], counting the frequency of concepts mentioned in the pre- and post-definition of Task 1 and Task 2 to determine the change in understanding. Then we thematically coded the transcript of participants' thinking aloud to discuss the behaviors that emerge when looking at a past design artefact with an AI Mentor. Thematic analysis was coded by multiple inter-coders. For RQ2: Understanding whether the CIAT interface was useful during brainstorming and how participants used it to proceed with ideation, we first analyzed the self-report survey in combination with behavioral logs. The survey questions, shared in Table 2, targets Creative Confidence and Growth [8] Usefulness of AI Mentor's feedback [27], and Intellectually Stimulated in novel ways [31]. Since perceived satisfaction with the final idea can be difficult to interpret, we then relate each participant's survey data to their interaction logs, which consist of the number of questions asked and counts of idea development

| ID | Gender | Design Expertise | AI Tools Usage for Ideation | Familiarity with Telepresence | Familiarity with Radical Atoms | In-person Relationship with Mentor |
|---|---|---|---|---|---|---|
| P1 | male | 10 years | sometimes | advanced | moderate | yes |
| P2 | male | 8 years | sometimes | moderate | basic | no |
| P3 | male | 5 years | often | advanced | advanced | yes |
| P4 | male | 3 years | sometimes | moderate | moderate | no |
| P5 | male | 8 years | sometimes | advanced | moderate | no |
| P6 | female | 8 years | often | advanced | moderate | no |
| P7 | female | 5 years | sometimes | moderate | basic | no |
| P8 | female | 2 years | rarely | moderate | basic | no |
| P9 | male | 2 years | often | moderate | basic | no |
| P10 | male | 5 years | often | moderate | basic | yes |
| P11 | male | 5 years | often | advanced | moderate | yes |
| P12 | female | 10 years | sometimes | advanced | advanced | yes |

**Table 3: Participant information for the study. For AI tool usage for ideation, rarely = tried 1–2 times, sometimes = occasional use, often = regularly integrated into ideation workflow. For familiarity with Telepresence and Radical Atoms, Advanced = understands core concepts and has done projects, Moderate = read papers and can name multiple projects, Basic = heard of the concept and can name 1–2 projects.**

and iteration, to gain a more holistic understanding of their brainstorming behavior. Lastly, we triangulate the quantitative findings with our rich, thematically analyzed qualitative data of participants' thinking out loud during task 2 and the semi-structured interview. Lastly, we investigated RQ3: The social and behavioral experience of co-ideation across time by analyzing the coded qualitative data from four tasks, with a focus on emotional experiences, and found interesting relationships to participants' familiarity with the mentor.

# 6 Findings

We present our findings on how participants comprehended sketchnotes with and without an AI Mentor (RQ1), in which ways participants interpret, apply, or reject the AI Mentor's generated input during brainstorming (RQ2), and what experiences and social behavior surface when participants ideate across time (RQ3).

## 6.1 Understanding Sketchnotes with AI Mentor.

We first summarize participants' definitions of Telepresence (without AI Mentor) and Radical Atoms (with AI Mentor) before and after examining the respective sketchnotes. We then summarize three strategies that participants apply to comprehend the sketchnote with the LLM tool compared to when using it without. We illustrate relevant findings with quotes.

Pre vs. Post Definition: The most frequent themes participants related to telepresence before looking at the sketchnote are Long-distance Communication (n=8) and Existing Tools References, e.g., Zoom (n=4). After examining the sketchnote, the themes shifted towards Sensory Presence (n=6), Emotional Connection (n=3), and Mental State of Being Present (n=3). Participants' understandings shifted from collaboration tools (e.g., Zoom, VR) to more embodied and experiential descriptions (e.g., "feeling present with another person"). In addition to reading the sketchnote, participants also reflected on the handwriting and thinking process of the human mentor. The most common themes associated with Radical Atoms prior to sketchnotes are Technical Shape-changing Material (n=6), Project References (n=3), and Material-based Interaction (n=3). After interacting with the sketchnotes and an AI Mentor, participants shifted from mechanical, ambiguous, or project-based understanding to a more design-related and specific understanding of Radical Atoms. The technical, physics-related aspect disappeared completely; instead, participants understood RA as a design framework, with Interaction Design Framework (n=6) as the most frequent theme. Material-based Interaction (n=3) remained a common definition, alongside a new theme, Environmentally Responsive Material (n=3). Compared to without an AI Mentor, this presents a deeper shift in understanding – from mainly using mechanical language to describe RA toward more conceptual and metaphorical descriptions.

Our qualitative data provided further details on how the participants interacted with the AI Mentor and what might have led to an enhanced comprehension. Participants viewing sketchnotes without AI Mentor learned by paraphrasing, while those in the AI condition validated and challenged their interpretation by asking questions. These questions targeted illegible handwriting, obtuse words, concept-specific words, and most interestingly, spatial and temporal reading order within the sketchnote. The generated responses were effective and allowed participants to address these basic hurdles upfront. This reduced the cognitive load significantly, keeping the sketchnotes interesting and allowing them to spend more time examining the content. We categorized 3 types of questions that participants asked the LLM, which would have remained unresolved without the AI Mentor.

**Improved Readability and Comprehension.** This type of question tackles illegible handwriting, abstruse words, and concept-specific terms. An intrinsic obstacle with sketchnotes that were mostly ignored and remained unresolved during Task 1. The AI Mentor, on the other hand, was effective in resolving these simple confusions. P7 asked, "What is 'I.O. Clay' (couldn't read Ill, thought it was IO) Illuminating clay! Oh, cool, clears that up". Clarifying concept-specific words was most helpful in improving the understanding. Oftentimes, during ideating, people like to come up with new terms to describe a new idea. These terms are not easily found online and can be difficult for generic agents to interpret without the appropriate context. With the AI Mentor, P2 was able to circle "sematic operator", and ask for further explanation and get a satisfying answer. P11 was curious about the Damascus knife scribble and asked AI Mentor to clarify why there is a sketch of a seemingly unrelated sword. Upon clarification that the knife serves as a metaphor for annealing layers of interaction into an object, such as a Damascus steel blade, the participant, who is experienced in traditional craftsmanship, immediately expressed a deeper reflection on the Radical Atoms concept. "I know Damascus steel. This is how they make the samurai swords. With many layers, it creates a wavy pattern of steel. Amazing, the RA concept got much clearer now". For the participants, using the AI Mentor to answer readability-related questions was perceived as "low cost", which encouraged them to ask questions that often might be perceived as "obvious". Later, P7 reflected that without an AI Mentor, "There is a high concept embedded in it somewhere, but I'm too lazy to really dig into it," whereas with an AI Mentor, "I know it's a big concept, so I just need to ask the right questions."

**Validating and Challenging one's previous Understanding.** When reading through the AI Mentor's response, participants would often compare the answer to the drawings and writings on the sketchnote. P4 reads the generated response "Environment participates like umbrellas, reacting to wind and rain.", then points to the sketchnote "I see this is an umbrella and it can change shape depending on the use case." Whenever there is a misalignment between the AI Mentor's response and the participant's perception, users feel challenged in their own understanding. In some cases, participants (P4, 10) disregarded the generated response and continued with their initial understanding; however, most participants returned to the sketchnote and tried to update their comprehension. P3 reads the "meaning tied to the object's role", then says:" Hmm. Not what I was thinking. I always thought RA is about new material, why is the Bot using everyday objects as examples for RA?" Cross-checking the generated response with what participants see on the sketchnote allowed them to critically engage and question rather than passively rely on generated responses. P6, who has moderate familiarity with RA, says, "The more I learn, I feel that the concept is more multifaceted than I originally thought. I feel like the paper I read only scratches the surface of the concept". Compared to without an AI Mentor, participants' comprehension relied more on guessing, with no means to validate or question their own assumptions. Our findings suggest that coupling an AI expert persona with genuine artifacts from the human mentor positions the users between blindly trusting a generated response and completely ignoring the AI Mentor's input. On the contrary, grounding generated text to an authentic piece of writing affords a more reflective and nuanced learning behavior with AI personas.//

**Supports comprehension of spatial relationships and temporal processes within Sketchnotes.** The task most often given to the AI Mentor was to help with the reading order and spatial relationships of writings and sketches. P11 started by circling a large section on the sketchnote and asking, "Summarize this section." They later explain, "It takes a lot of energy trying to understand all the handwriting." Since learning is already challenging, participants like to offload cognitively overwhelming tasks, such as structuring the reading order of the sketchnote or understanding the context of the writing, to the AI Mentor. Similarly, P7's first question is "What is the purpose of the word 'environment' here, or what is the context?" Establishing a roadmap of how each segment of a sketchnote relates to one another helps participants make sense of the seemingly scattered information and create a conducive learning environment for further exploration. Participants more familiar with RA projects attempted to identify the beginnings of published projects they were familiar with, to better understand the sketchnote. P6 "Has any of these concepts been turned into projects? If so, which one?" Participant later shared, "I think it [knowing the final outcome] makes me more curious about the concept. (...) It's really fun to think about all those iterations, all those ideas that led him from here to all the products he's been doing afterwards." The temporal relationship between sketchnotes and final projects adds another layer for participants to deepen their natural learning and comprehension. P10 reflects "When I see a piece of art or project, I always wonder, where did it come from? It's very interesting for me to trace it. I think the sketchnote with the Bot helps to answer that question very effectively." In summary, with the AI Mentor, participants were able to ask questions that would often remain unresolved when reading sketchnotes, including clarification of readability, request critiques to validate or challenge their understanding, and questions about the spatial relationships within a sketchnote. The back and forth between sketchnote and generated responses created opportunities for deeper inquiry and reflection on the the participants' own interpretations of the concept.

## 6.2 Ideation Strategy with AI Mentors

We first provide an overview of how participants perceived the brainstorming session and its outcome, based on their self-report survey. Then we relate the survey to interaction logs captured automatically by the CIAT system, such as the number of times a participant asked a question. Lastly, we triangulate the quantitative findings with our rich qualitative data and summarize 3 ways participants interpreted AI Mentor's feedback to further their ideation.

*Quantitative Analysis*

**Self-Report Measures.** Participants expressed consistently high levels of creative engagement. They felt capable (Mdn = 6/7) and motivated (Mdn = 6/7) to generate Radical Atoms–aligned ideas. They also perceived that their ideas progressed over the course of the session (Mdn = 6/7). Ratings were highest for exploring novel directions, where participants strongly agreed with branching into new ideas (Mdn = 7/7). The widest variability in responses appeared in participants' self-reported engagement (range = 3–7). While many reported high engagements, others expressed substantially lower engagement. We hypothesize that this spread may reflect differences in trust toward the AI Mentor, with less trusting participants engaging less deep with its suggestions.

**Interaction Logs Analysis.** The analysis of interaction logs revealed more distinct differences in how participants engaged with the AI Mentor during brainstorming. Some participants emphasized guidance-seeking behaviors, asking frequent questions during ideation, e.g., P4 and P12, with 16 questions each. Others

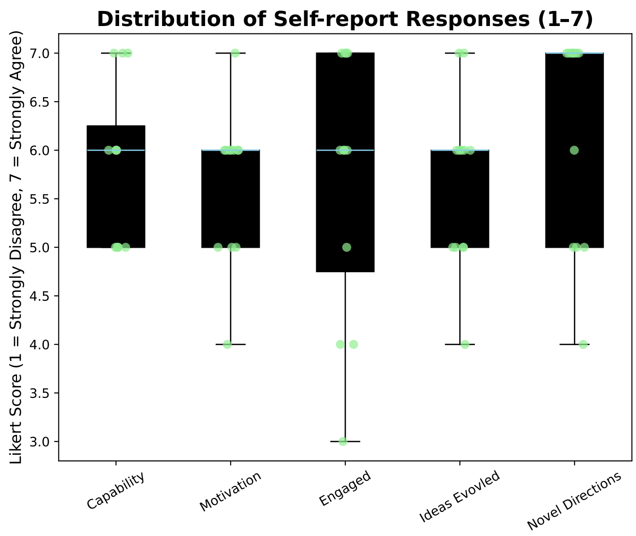


**Figure 7: Self-reported data scores**

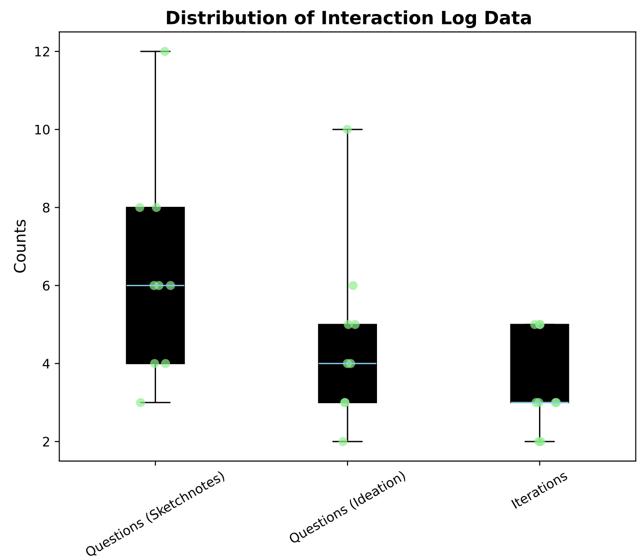


**Figure 8: Distribution of times of interaction automatically captured by Co-ideation Interface system log.**

engaged primarily in feedback-driven iteration, repeatedly refining and revising their ideas through iteration, refining ideas multiple times, e.g., P5 and P7, each with 5 iterations.

**Combining.** self-report measures with interaction logs analysis revealed that participants who gave high novelty scores for their final idea tend to ask more questions and propel their ideas into unexpected conceptual spaces. P12 began with creating a bento box that shape-shifts to fit its contents, but ended with a sketch of a bento box that turns throwing away leftovers into a mourning ritual. By contrast, participants who reported their ideas progressed over time had higher iteration counts. They produced more layered and progressively refined concepts. P7 started with something out of fabric to an actuated piece of textile to an actuated textile that would respond to the user's emotions and ended with a memory and experience recreation quilt.

*Qualitative Data Analysis*

In relation to the previously discussed quantitative findings, we report the strategies participants used to utilize AI Mentor-generated feedback to advance their idea, foster creative confidence, and maintain ownership over the outcome.

**AI Mentor as a Mirror for Thoughts.** Articulating one's own thought is a key outcome of ideation [5]. Therefore, when participants shared their half-baked idea with the AI Mentor and asked for feedback, many interpreted AI's response to structure their own ideas and help themselves proceed. P8 "In my ideation sessions, I usually don't think about what verb I want to use or which qualities I really care about. (...), but once you see these questions listed out, it becomes very clear which direction I'm interested in. P9 "I first went in with objects in mind, and then I ideated with the bot. It then ranked my thoughts. So that order helped me order my own ideas..." AI expert personas are specific and, in this case, paraphrase participants' ideas using HCI vocabulary while also prompting them to consider novel directions that are grounded in mentor's published work. This is useful for participants when they want diverse ideas, but for a specific audience. P12 "I was amused. There were many things that came from the bot that I would not have thought of, because I'm not the mentor." By using the AI Mentor as a proxy to clarify ideas, rehearse explanations, and learn about domain-specific jargon, participants' self-reports reflect a high level of felt capability, confidence, and motivation to generate RA ideas.

**Nuanced Lens to Adapt Generated Feedback.** We observed that participants had 3 different ways of categorizing AI Mentor's feedback. Direct Feedback refers to instances where participants identified AI Mentor's feedback as directly actionable for their brainstorming and were willing to take the suggestion. This behavior was mostly triggered by generated responses that included a question, such as "Choose one (from a list of generated options)" or "What is your core interaction?" P4, for example evaluates the different interaction metaphors AI Mentor provided," Boxes is a cool one! I kinda like boxes. Candle not really, I don't see how Candle gonna work. Clock… Clock is a cool one. Mirror is overdone. Bell's unrelated. Let's do clock." and then moved on to a more focused brainstorming. However, adopting feedback does not necessarily mean immediately incorporating it into their current idea.

- **Long-term Feedback** refers to AI Mentor's responses that were understood as pieces of advice for future work and triggered reflection on their usual ideation process. P5 was asked by the AI Mentor to focus on one core interaction and responded, "I want to do both. But If I had to choose, probably the ergonomic (grip) one is more important. But (I prefer) the beauty (that) lies in the cultural continuity of the tool. I don't want to choose. I do I understand, though, why I need to choose, because this is too vague of a concept." This is interesting, because with general LLMs, P2, 3, 5, 8, 11 reflected, they would only pick out the responses that are applicable to them at that moment and ignore the rest. However. P8 compares "[this experience] is influential because it gave me an idea to hold. I don't really think about the human body's relationship with objects. I didn't really consider that aspect before. That's an interesting new direction I will be thinking more about and try to consider in my future ideas."

- **Personal Preference Feedback.** Unique to the AI Mentor compared is that participants were able to interpret the AI Mentor responses as the bot's preference or style rather than a direct critique of their idea. P4 remarks that they think the AI Mentor is really good for ideating on tactile ideas. "It continuously talks about affordances such as rewinding and tightening the clock gesture." This suggests that participants can first determine whether the feedback is generalizable or simply a preference of the AI Mentor due to its training, before adapting it. The custom conversational agent provided the participant with a more nuanced lens to interpret the generated responses.

Overall, we observe that AI Mentor's feedback helps participants to progress with their idea. While our exploratory study did not formally compare AI Mentor's feedback to that of a general LLM, participants naturally recalled their existing experience of brainstorming with ChatGPT. Our analysis of those reflections shows that participants treat general LLMs more like a tool to increase productivity, while ideating with the AI Mentor, participants engaged more deeply and tried to learn or draw lessons from the generated responses.

**Entering the Mind of the Author of the Sketchnote.** Similar to Task 1, during ideation, participants also frequently compared the AI Mentor's feedback with what they think the actual author might say. This not only allowed them to easily identify AI slop, P5 "there are like 3 (feedbacks) that are a little bit random and kind of bad feedback, and then, like, one or two that are good mentor name feedback." but furthermore to reflect on what the author of the sketchnote might think of their idea. P1 " "Poetic metaphor, but turned into an interaction. . . ", (. . . ) Good feedback. I would imagine this from *human mentor's name*". Some even described ideation as a 3 player effort (participant, AI Mentor and human mentor), stating that they were able to "channel the real Mentor" 3 or 4 times, where as in their usual ideation sessions, they would only be able to do so limitedly. This effect was especially strong, when the AI Mentor referenced and showed a picture of a sketchnote during ideation. This triggered curiosity and a feeling of picking up ideas across time. P11 shares in the interview that it felt "like "enter(ing) into their mind, when they were coming up with that moment of creation." which is intellectually stimulating and very fun.

Overall, we found that co-ideation interface supports user's personalized ideation strategies from offering structure, variation of types of feedback to provoking intellectual engagement across time. Moreover, we found that all participants played the dominant role in the partnership, preferring to maintain control over the design direction. This represents a positive outcome when compared to traditional human-human ideation, where it can be difficult to balance the contributions of each collaborator. The question of agency and other social experiences will be discuss in the next section.//

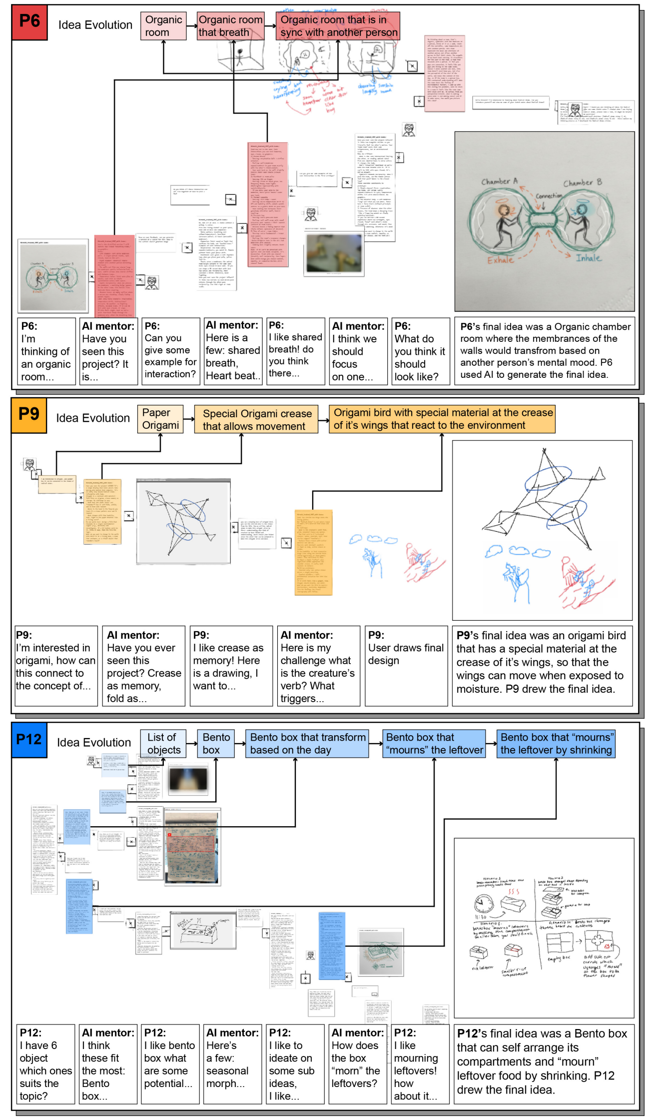


**Figure 9: Three examples of ideation evolution using the Co-Ideation Interface**

## 6.3 Experiences and Social Behavior of Co-Ideation Across Time

In this section, we report on participants' experiences and social behaviors with a system that combines legacy design sketchnotes with an interactive AI Mentor. We summarize the AI Mentor's perception and why it might have led to a strengthened sense of creative ownership. Then we report on participants' feelings of reduced social pressure when interacting with the AI Mentor. Lastly, we reflect on the learning how learning across time differs based on participants' existing relationship to the human mentor.

**Strengthened Feeling of Ownership.** In our analysis, we found that the feeling of ownership over the final idea might be linked with participants' perception of the AI Mentor as an authoritative and inspirational voice. P1 states that they are glad the AI Mentor is trained on one person's knowledge. While it makes the responses very specific, it gives credibility too: "he has a more authoritative voice. Whereas regular ChatGPT is just a blend of online people's opinions." Since the agent learned from how the real author frames his ideas, the AI Mentor is more inspirational. "I think a lot of my ideas lack poeticism, and this is a nice solution... Where if you tell a generic AI to be poetic, it most likely goes in the wrong poetic direction," commented P12. The sketchnotes played a vital role in establishing the credibility of the AI Mentor. P9 articulated that for them, seeing the sketchnote was the most important step. "There is a cognitive shift that happens when I see the sketchnotes; it gives the conversation a backstory." It allows participants to transition from simply following or rejecting the AI's suggestions to becoming interested in RA and willing to contribute their own input, which may ultimately lead to a sense of ownership over the final outcome. While the AI Mentor is seen as a figure of authority, the participants never confused the LLM persona with the real person. P4 shares, "From chatting through the sketches with it, I feel it does have enough knowledge. This is just a clone, but it's a good enough clone, and that's enough." This clear understanding of the AI Mentor as a tool, rather than a replacement for a real person, helps participants maintain creative freedom and agency throughout the experience. In some cases, this distinction made it even easier for participants to feel ownership over the idea. P11 is often hesitant to share ideas too early with their mentors, worried that they might feel pressured to follow the mentor's feedback and gradually lose interest in their own idea. For this experience, however, he reflected that he could see himself involving AI Mentor early on to receive the desired domain-specific feedback. "I'm 100% in control and did not have to do what it told me to. I still feel ownership over the final idea that came out, even though it was built on top of their work," P11 reflected.

**Free of Social Pressure.** All 12 participants mentioned feeling less pressured by social expectations when interacting with the AI Mentor compared to the human mentor or even their peers. Some felt a layer of respect for more senior researchers. P10 "I know I'm only wasting the bot's time. If I come up with a bad idea, there's no pressure. That helped." Others described how they would appreciate input when their "idea is still being formed"; however, often in fear of judgment or lack of words, they would not seek feedback from an expert. P6 says, "I probably would never tell it to the mentor in person, because the idea is vague and hard to describe. But with the Bot, I just try. Maybe I can refine my thinking, refine my question, step by step". Interestingly, P12, who has received in-person feedback from the human mentor before, found new appreciation for the previous conversations with the mentor after interacting with the bot, saying, "I think that there is a poeticism to the framing of the mentor's work. (...) A lot of my ideas lack poeticism. This is a nice solution for me to understand that poetic notion. Where with normal AI, it usually goes in the wrong poetic direction."

**Learning Across Time.** While not part of our initial research question, we found that using a past sketchnote with traces of finished projects that participants might be familiar with today has a significant effect on the emotional experience. We find that especially participants with an existing relationship to the human mentor resonated more with the sketchnotes. Participants who were only familiar with the mentor's work primarily focused on learning the content. While people with knowledge of publications or who had in-person conversations with the actual mentor connected the sketchnotes to a bigger picture and were able to profit more from cross-generational knowledge sharing. They could identify sketchnote sketches as the beginning of future projects and described the experience of co-ideating across time as "not just look(ing) into the past, but projecting the future (P6)." Some were surprised to find that seeds of concepts being developed in 2025 had as far back as 2015. P11 remarks, "I also felt a lot of today's mentors' interest coming from it, even though I know it wasn't the main focus back then." Through relating the past and projecting into the future, participants gained insight into the struggles of brainstorming for a now-established senior researcher. "I saw the ideation. (. . . ) I saw the work that the mentor put in. And the work gave me understanding of his intention, and the intention resonates with my own intention of why I'm working on my current projects." P2 comments. Others found the human mentor to be more relatable and felt their own goals as a creative person is more achievable. P6 shares, "His sketches look quite familiar to how we ideate. I think it brings the distance closer, seeing he also did these things before. It feels like, if I have an idea, I can also build my own thinking."

**Interface Design Feedback.** Overall, participants appreciated the CIAT interface design. One positive feature is that the ideation process is centralized, eliminating the need to switch interfaces to find references. P9 "This is a space where I can work on something undistractedly. I can ideate in the most intuitive way, which is sketching a prototype, ideas, and getting responses in the same space". Furthermore, users appreciated the node-based system, which enabled them to connect ideas and visualize how their ideas evolved from the original inspiration to the final feedback. There are also many areas for improvement. Several users critiqued the lengthiness, wordiness, and sometimes randomness of the AI Mentor's language. One also pointed out that while they enjoyed looking at the sketchnote, they found it cumbersome to have to probe the AI Mentor for an explanation of the sketchnote and would much appreciate a clean, abstracted outline in addition to the original.

## 7 Discussion

Our paper studies how an AI Mentor can support learning from sketchnotes, facilitate co-creation, and shape interpersonal and reflective experiences. Across all three research questions, we highlight how AI can scaffold interpretation and interpersonal dynamics while also raising new ethical questions about representation and access.

## 7.1 Spatial-temporal Interpretation of Sketchnotes

Participants using the AI Mentor exhibited noticeably higher engagement, curiosity, and willingness to persist with complex sketchnotes. This addresses a long-standing barrier identified by prior work [15, 42] that sketchnotes are often difficult to interpret and require specialized visual literacy skills. The AI Mentor acted as a bridge, reducing the cognitive overhead needed to "enter" a sketchnote. We observed three core strategies through which the AI Mentor augmented comprehension. 1) Improving Readability and Vocabulary. 2) Validating or Challenging Interpretations. 3) Unpacking Temporal and Spatial Structures. We argue that this third capability, spatial–temporal interpretation, is particularly interesting and positions the AI Mentor as an interpreter of visual narrative. While it is currently helping students comprehend the order and spatial relationships of sketchnote elements in context, we argue that it can be further developed in the future by assessing the quality of handwriting or strokes to differentiate between emerging ideas and task management. This could allow for a new reading order across collections of sketchnotes to enhance learning without flattening the artistic qualities. This interpretive framing can extend to other embodied knowledge or artistic formats, such as annotating dance choreography or even artworks, like Jackson Pollock's paintings.

**Design Guideline 1: Support Structural Interpretation, Not Just Content Explanation.** AI Mentors for understanding design legacies should be designed to recognize visual narrative structure such as temporal sequences, spatial groupings, emphasis, and flow, and communicate these structures back to users. Rather than merely explaining symbols, an AI Mentor should help learners traverse the intended (or alternative) reading paths of complex visual artifacts.

## 7.2 Physical Presence of Sketchnotes

While our findings align with prior research showing that AI can support ideation, reasoning, and expand creative search [12, 26, 28, 51], we identify a novel effect: the physical presence of the mentor's sketchnotes lent the AI Mentor stronger perceived authority and "intellectual presence." Participants felt that the agent was not generating ideas in the abstract but responding to the human mentor's personal artifacts, history, and perspective. In this sense, the sketchnote becomes a prop that grounds the LLM's feedback, anchoring it in a person's embodied trace [22]. This tangibility stands in contrast to standalone creativity tools, which often feel detached or decontextualized [44]. A future question that emerges is whether handwritten sketchnotes carry different authority or emotional resonance than digitally cleaned-up versions, such as a polished diagram. Handwriting may surface personality, giving the AI more authentic material from which to infer intent.

**Design Guideline 2: Ground AI Mentor's Feedback in Traceable Artifacts.** AI Mentors should incorporate hand-created, tangible, or genuine materials, such as sketchnotes or messy drafts during co-ideation, to provide contextualized, personally relevant support for brainstorming. The system should treat these artifacts as "anchors" to guide the conversation and provide users with more means to determine whether an impersonating LLM's response is trustworthy or not.

## 7.3 Personal Distance Affects AI Mentor Experience

Participants reported that co-creating with the AI Mentor felt "free of social pressure" and made it easier for some to admit confusion, echoing patterns described by Remirrorfugue [53] and related work on working with representations of mentors. However, the experience differed distinctively between participants who knew of the mentor in person and those who only knew of their work. Those who felt no relationship focused on learning the concepts, while those who had a prior relationship used the AI as a proxy to reflect on their professional relationship with the mentor. Some even found this experience useful in reinterpretation of previously given feedback and better understand the in-person feedback. These findings raise deeper design questions: How might we train different AI Mentors for different audiences? Can an AI Mentor adapt to learners at various career stages differently? How can using the AI Mentor enrich the in-person relationship? Can computers enrich intergenerational knowledge sharing and connect people with similar interests even beyond our lifespans [24]?

**Design Guideline 3: Allow Users to Calibrate Distance With the AI Mentor.** One can think more nuanced when creating AI Mentors for intergenerational knowledge transfer and co-ideation. From different personalities, ranging from highly supportive and exploratory to concise and analytical, so users can select the degree of social presence they need. Furthermore, there could be different versions of the same AI Mentor from different time periods. Often times human mentor's interests invovle. This does not mean their previous interest is irrelevant, but it is more difficult to answer student's questions of a topic one thought about 10 years ago. In this way, the systems could also communicate the limits of these personas.

## 7.4 Ethical Considerations

If expert-perspective LLMs are used for intergenerational knowledge transfer, we must confront questions of authority, representation, mentor involvement, transparency of limitations, and credit/ownership. Who gets to serve as the "voice" of a mentor? How do we communicate what the AI can and cannot reliably do? And how should authorship be attributed when the AI extends or transforms an idea? Importantly, our intention is not for AI Mentors to replace real mentorship, but to enrich and scaffold conversations that continue to happen in person. These questions frame the ethical stakes of encoding human expertise into AI systems.

First, we observed that users appreciated the AI Mentor's poetic tone, even while recognizing that these were not the original words of the human mentor. If the AI deviates too significantly, such as presenting ideas that the mentor never endorsed, it may risk undermining users' trust and misrepresentation. Given these concerns, a first-person narrative ("I am your mentor") may blur boundaries and overstate what the AI can legitimately claim. A third-person framing ("This is an AI that draws on the mentor's materials") may better preserve clarity and respect for the human mentor. Therefore,

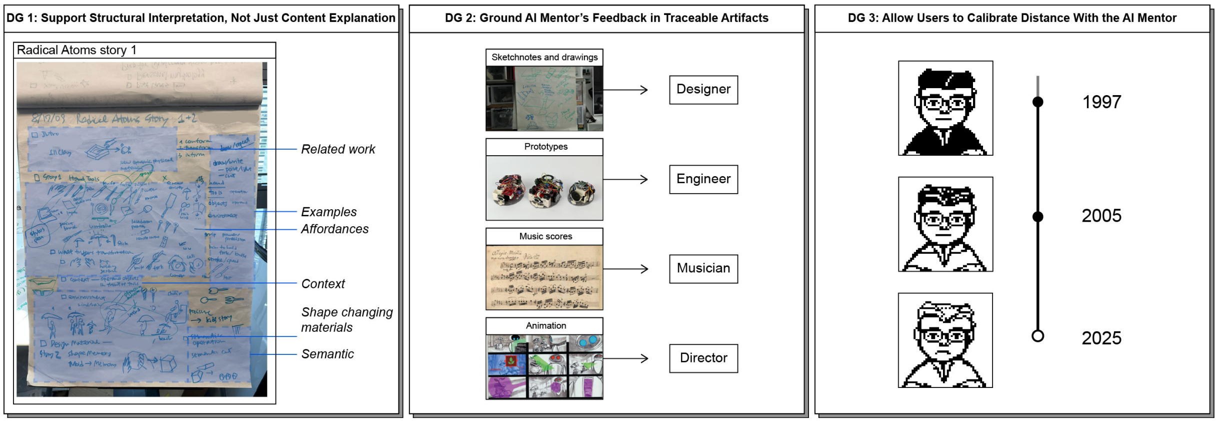


**Figure 10: Design Guidelines**

defining the appropriate boundary for AI-driven augmentation is a critical design consideration. This boundary might also differ depending on the context, for instance, when designing for the legacy of a mentor who has passed away versus one who is still living.

Second, the involvement of the author of the sketchnote is crucial not only for accuracy but for transparency. Our system's flexibility supports both reflection and exploration; however, without clear signals indicating when the AI was channeling publication-derived content versus improvising new interpretations, users sometimes experienced uncertainty or cognitive overload. For example, participants might delve deeply into a single concept generated by the AI Mentor, which is not the main focus of the author. Further exploration into when and how one should make the AI Mentor's decision-making visible, defining which parts of the system reflect their intentional contributions, and clarifying when the AI is generating novel content. Ultimately, ownership and attribution remain crucial as sketchnotes evolve into training materials for AI. Many creators hesitate to share work-in-progress material due to concerns about judgment [45] or loss of control. If their materials are used to train or shape AI Mentors, transparent consent processes and clear attribution practices become essential.

## 8 Limitations

We acknowledge several limitations in both our methodology and system design. Regarding methodology, we did not counterbalance the order of interpreting the sketchnote between conditions with and without AI. This is a possible limitation that could affect the interpretation of the quantitative results. While we attempted to mitigate potential order effects, more extensive studies are necessary to confirm whether the findings remain consistent across different concepts. In terms of system design, due to technical usability issues, 3 participants were unable to directly control the mouse and keyboard and instead asked the research facilitator to do so. This may have influenced participants' learning experience and a portion of the data collected; however, most participants interacted with the system as intended. Additionally, the AI Mentor's feedback was sometimes lengthy, making it difficult for users to read.

## 9 Conclusion and Future Work

This paper demonstrates the potential of AI to revitalize and reinterpret legacy design sketches, enabling meaningful co-ideation across academic generations. Our system unlocks tacit knowledge through an AI Mentors and AI-augmented Knowledge Objects created based on the theses and publications of human mentors. The exploratory evaluation with 12 participants confirms that the system stimulates learning, sparks novel design directions, enhances cross-referencing, and evokes a sense of learning across time. While people often think of design archives as static repositories, we revitalize legacy design sketchnotes as AI-augmented knowledge objects to become active participants in collaboration. Our future work includes continuing to iterate on the system design, expanding into different knowledge domains, such as architecture and engineering, training multiple AI Mentors who can provide insights from various perspectives, and allowing for multiple users to facilitate group ideation. Looking forward, this research opens new avenues for AI-supported intergenerational knowledge sharing and design education.

## Acknowledgments

We thank Prof. Min Fan without whose expertise this paper would have never made it into the first round and we thank of course the TeleAbsence vision.

## References

[1] [n. d.]. AI-Augmented Brainwriting - Google Suche. https://dl.acm.org/doi/10.1145/3613904.3642414
[2] [n. d.]. Industry-scale Knowledge Graphs: Lessons and Challenges. https://research.google/pubs/industry-scale-knowledge-graphs-lessons-and-challenges/
[3] 2025. (PDF) Development of a Lightweight AI-Powered Personal Agent Using OpenAI's GPT-4 API. https://www.researchgate.net/publication/391702034_Development_of_a_Lightweight_AI-Powered_Personal_Agent_Using_OpenAI's_GPT-4_API
[4] Amin Beheshti. 2024. Natural Language-Oriented Programming (NLOP): Towards Democratizing Software Creation. https://doi.org/10.48550/arXiv.2406.05409 arXiv:2406.05409 [cs].
[5] Bill Buxton. 2007. *Sketching User Experiences: Getting the Design Right and the Right Design.* Morgan Kaufmann, Amsterdam Boston.
[6] 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. ACM, Delft Netherlands, 1–11. https://doi.org/10.1145/3582269.3615596
[7] Valerie J. Caracelli and Jennifer C. Greene. 1993. Data Analysis Strategies for Mixed-Method Evaluation Designs. *Educational Evaluation and Policy Analysis* 15, 2 (June 1993), 195–207. https://doi.org/10.3102/01623737015002195
[8] Gilad Chen, Stanley M. Gully, and Dov Eden. 2001. Validation of a New General Self-Efficacy Scale. *Organizational Research Methods* 4, 1 (Jan. 2001), 62–83. https://doi.org/10.1177/109442810141004
[9] Luigina Ciolfi, Areti Damala, Eva Hornecker, Monika Lechner, Laura Maye, and Daniela Petrelli. 2015. Cultural heritage communities: technologies and challenges. In *Proceedings of the 7th International Conference on Communities and Technologies (C&T '15)*. Association for Computing Machinery, New York, NY, USA, 149–152. https://doi.org/10.1145/2768545.2768560
[10] Jeff Conklin and L Begeman. [n. d.]. glBIS: A Hypertext Tool for Exploratory Policy Discussion. *ACM Transactions on Office Information Systems* 6, 4 ([n. d.]).
[11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. https://doi.org/10.18653/v1/N19-1423
[12] Giulia Di Fede, Davide Rocchesso, Steven P. Dow, and Salvatore Andolina. 2022. The Idea Machine: LLM-based Expansion, Rewriting, Combination, and Suggestion of Ideas. In *Proceedings of the 14th Conference on Creativity and Cognition (C&C '22)*. Association for Computing Machinery, New York, NY, USA, 623–627. https://doi.org/10.1145/3527927.3535197
[13] Robert Dimeo. 2016. Sketchnoting: an analog skill in the digital age. *SIGCAS Comput. Soc.* 46, 3 (Dec. 2016), 9–16. https://doi.org/10.1145/3024949.3024951
[14] Claudia Eckert and Martin Stacey. 2000. Sources of inspiration: a language of design. *Design Studies* 21, 5 (Sept. 2000), 523–538. https://doi.org/10.1016/S0142-694X(00)00022-3
[15] Marina Fernández Camporro and Nicolai Marquardt. 2020. Live Sketchnoting Across Platforms: Exploring the Potential and Limitations of Analogue and Digital Tools. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–12. https://doi.org/10.1145/3313831.3376192
[16] Bill Gaver and John Bowers. 2012. Annotated portfolios. *Interactions* 19, 4 (July 2012), 40–49. https://doi.org/10.1145/2212877.2212889
[17] Katrin Glinka and Claudia Müller-Birn. 2023. Critical-Reflective Human-AI Collaboration: Exploring Computational Tools for Art Historical Image Retrieval. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2 (Oct. 2023), 263:1–263:33. https://doi.org/10.1145/3610054
[18] Maria Gäde, Marijn Koolen, Mark Hall, Toine Bogers, and Vivien Petras. 2021. A Manifesto on Resource Re-Use in Interactive Information Retrieval. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*. ACM, Canberra ACT Australia, 141–149. https://doi.org/10.1145/3406522.3446056
[19] Zhiting He, Jiayi Su, Li Chen, Tianqi Wang, and Ray Lc. 2025. 'I Recall the Past': Exploring How People Collaborate with Generative AI to Create Cultural Heritage Narratives. *Proc. ACM Hum.-Comput. Interact.* 9, 2 (May 2025), CSCW108:1–CSCW108:30. https://doi.org/10.1145/3711006
[20] Scarlett R. Herring, Chia-Chen Chang, Jesse Krantzler, and Brian P. Bailey. 2009. Getting inspired!: understanding how and why examples are used in creative design practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Boston MA USA, 87–96. https://doi.org/10.1145/1518701.1518717
[21] Angel Hsing-Chi Hwang and Andrea Stevenson Won. 2021. IdeaBot: Investigating Social Facilitation in Human-Machine Team Creativity. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–16. https://doi.org/10.1145/3411764.3445270
[22] Hiroshi Ishii. [n. d.]. "The Last Farewell" Traces of Physical Presence. ([n. d.]).
[23] Hiroshi Ishii, Dávid Lakatos, Leonardo Bonanni, and Jean-Baptiste Labrune. 2012. Radical atoms: beyond tangible bits, toward transformable materials. *Interactions* 19, 1 (Jan. 2012), 38–51. https://doi.org/10.1145/2065327.2065337
[24] Hiroshi Ishii, Daniel Pillis, Pat Pataranutaporn, Xiao Xiao, Hayoun Noh, Lucy Li, Alaa Algargoosh, and Jean-Baptiste Labrune. 2025. TeleAbsence: A Vision of Past and Afterlife Telepresence. *MIT Press* (2025). https://doi.org/10.1162/PRES_a_00441 Accepted: 2025-03-12T15:14:07Z.
[25] David G. Jansson and Steven M. Smith. 1991. Design fixation. *Design Studies* 12, 1 (Jan. 1991), 3–11. https://doi.org/10.1016/0142-694X(91)90003-F
[26] Ji-Youn Jung, Devansh Saxena, Minjung Park, Jini Kim, Jodi Forlizzi, Ken Holstein, and John Zimmerman. 2025. Making the Right Thing: Bridging HCI and Responsible AI in Early-Stage AI Concept Selection. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. ACM, Madeira Portugal, 2992–3012. https://doi.org/10.1145/3715336.3735745
[27] Pegah Karimi, Kazjon Grace, Mary Lou Maher, and Nicholas Davis. 2018. Evaluating Creativity in Computational Co-Creative Systems. https://doi.org/10.48550/arXiv.1807.09886 arXiv:1807.09886 [cs].
[28] Yewon Kim, Sung-Ju Lee, and Chris Donahue. 2025. Amuse: Human-AI Collaborative Songwriting with Multimodal Inspirations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–28. https://doi.org/10.1145/3706598.3713818
[29] Klaus Krippendorff. 2013. *Content Analysis: An Introduction to Its Methodology*. SAGE. Google-Books-ID: s_yqFXnGgjQC.
[30] Caitlin Marie Lancaster, Wen Duan, Rohit Mallick, and Nathan J. McNeese. 2025. Human-Centered Team Training for Human-AI Teams: From Training with AI Tools to Training for AI Teammates. *Proc. ACM Hum.-Comput. Interact.* 9, 2 (May 2025), CSCW100:1–CSCW100:38. https://doi.org/10.1145/3710998
[31] Tomas Lawton, Kazjon Grace, and Francisco J Ibarrola. 2023. When is a Tool a Tool? User Perceptions of System Agency in Human–AI Co-Creative Drawing. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 1978–1996. https://doi.org/10.1145/3563657.3595977
[32] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81. https://aclanthology.org/W04-1013/
[33] Yimeng Liu. 2024. Interaction Design for Human-AI Choreography Co-creation. https://doi.org/10.48550/arXiv.2405.03999 arXiv:2405.03999 [cs].
[34] Yiren Liu, Pranav Sharma, Mehul Oswal, Haijun Xia, and Yun Huang. 2025. PersonaFlow: Designing LLM-Simulated Expert Perspectives for Enhanced Research Ideation. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 506–534. https://doi.org/10.1145/3715336.3735789
[35] Wenhan Lyu, Yimeng Wang, Yifan Sun, and Yixuan Zhang. 2025. Will Your Next Pair Programming Partner Be Human? An Empirical Evaluation of Generative AI as a Collaborative Teammate in a Semester-Long Classroom Setting. In *Proceedings of the Twelfth ACM Conference on Learning @ Scale (L@S '25)*. Association for Computing Machinery, New York, NY, USA, 83–94. https://doi.org/10.1145/3698205.3729544
[36] Iain Macdonald. 2014. Digital And Paper Sketchbooks: Learning New Ways Of Drawing And Designing. *Tracey: drawing and visualisation research* (Feb. 2014). http://www.lboro.ac.uk/microsites/sota/tracey/journal/sbook/macdonald.html
[37] Sarnoff Mednick. 1962. The associative basis of the creative process. *Psychological Review* 69, 3 (1962), 220–232. https://doi.org/10.1037/h0048850 Place: US.
[38] Marvin Minsky. [n. d.]. Telepresence. https://philpapers.org/rec/MINT
[39] Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-Writing Screenplays and Theatre Scripts with Language Models: Evaluation by Industry Professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–34. https://doi.org/10.1145/3544548.3581225
[40] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics (ACL '02)*. Association for Computational Linguistics, USA, 311–318. https://doi.org/10.3115/1073083.1073135
[41] Mike Rohde. 2013. *Sketchnote Handbook, The: the illustrated guide to visual note taking*. Peachpit Press, San Francisco, CA.
[42] Katharina Scheiter, Katrin Schleinschok, and Shaaron Ainsworth. 2017. Why Sketching May Aid Learning From Science Texts: Contrasting Sketching With Written Explanations. *Topics in Cognitive Science* 9, 4 (2017), 866–882. https://doi.org/10.1111/tops.12261 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/tops.12261.
[43] 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. https://doi.org/10.1145/3613904.3642414
[44] Joongi Shin, Anna Polyanskaya, Andrés Lucero, and Antti Oulasvirta. 2025. No Evidence for LLMs Being Useful in Problem Reframing. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–25. https://doi.org/10.1145/3706598.3713273
[45] Ben Shneiderman. 2007. Creativity support tools: accelerating discovery and innovation. *Commun. ACM* 50, 12 (Dec. 2007), 20–32. https://doi.org/10.1145/1323688.1323689
[46] Pao Siangliulue, Joel Chan, Krzysztof Z. Gajos, and Steven P. Dow. 2015. Providing Timely Examples Improves the Quantity and Quality of Generated Ideas. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*. ACM, Glasgow United Kingdom, 83–92. https://doi.org/10.1145/2757226.2757230
[47] Sarah Sterman, Molly Jane Nicholas, Janaki Vivrekar, Jessica R Mindel, and Eric Paulos. 2023. Kaleidoscope: A Reflective Documentation Tool for a User Interface Design Course. 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. https://doi.org/10.1145/3544548.3581255
[48] Miriam Sturdee, Makayla Lewis, and Nicolai Marquardt. 2018. SketchBlog #1: the rise and rise of the sketchnote. *interactions* 25, 6 (Oct. 2018), 6–8. https://doi.org/10.1145/3281661

[49] Samangi Wadinambiarachchi, Ryan M. Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The Effects of Generative AI on Design Fixation and Divergent Thinking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–18. https://doi.org/10.1145/3613904.3642919

[50] Shaun Wallace, Brendan Le, Luis A. Leiva, Aman Haq, Ari Kintisch, Gabrielle Bufrem, Linda Chang, and Jeff Huang. 2020. Sketchy: Drawing Inspiration from the Crowd. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (Oct. 2020), 1–27. https://doi.org/10.1145/3415243

[51] 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. https://doi.org/10.1145/3706598.3714148

[52] T. B. Ward. 1994. Structured Imagination: the Role of Category Structure in Exemplar Generation. *Cognitive Psychology* 27, 1 (Aug. 1994), 1–40. https://doi.org/10.1006/cogp.1994.1010

[53] Xiao Xiao, Hayoun Noh, Adrien Lefevre, Lucy Li, Holly McKee, Alaa Algargoosh, and Hiroshi Ishii. 2025. ReMirrorFugue: Examining the Emotional Experience of Presence and (Illusory) Communications Across Time. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–26. https://doi.org/10.1145/3706598.3713328

[54] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating Text Generation with BERT. https://doi.org/10.48550/arXiv.1904.09675 arXiv:1904.09675 [cs].

[55] Rebecca Zheng, Marina Fernández Camporro, Hugo Romat, Nathalie Henry Riche, Benjamin Bach, Fanny Chevalier, Ken Hinckley, and Nicolai Marquardt. 2021. Sketchnote Components, Design Space Dimensions, and Strategies for Effective Visual Note Taking. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3411764.3445508