Tangible Co-Ideation: Designing Embodied Prompting for Creative Thinking with Large Language Models

Yuqing Lucy Li*
MIT Media Lab
Cambridge, USA
lucy_li@media.mit.edu

Quincy Kuang*
MIT Media Lab
Cambridge, USA
quincyku@media.mit.edu

Paul-Peter Arslan*
DeVinci Institute of Technology
Paris, France
paul-peter.arslan@devinci.fr

Yuhan Wang*
MIT Media Lab
Cambridge, USA
wyuhan@mit.edu

Xiao Xiao
Devinci Institute of Technology
Paris, France
xiao.xiao@devinci.fr

Hiroshi Ishii
MIT Media Lab
Cambridge, USA
ishii@media.mit.edu

ABSTRACT

While typing is the most common modality to instruct an LLM, there are applications where text can limit our interaction with AI. Especially during design, individuals might rely on visual thinking, imagination, and intuition, all of which do not follow a sequential chain of thought and can be challenging to express through sentences. Can tangible interfaces use embodiment to structure and craft prompts that challenge the linear thinking encoded into conversational LLM interfaces? We designed 3 object prompts that leverage tangibility to enable LLM interaction while mitigating information overload. The first group of objects, AI Expert Personas embodied as figurines, can fine-tune models depending on spatial position and distance. Second is a rotary lighthouse tower that searches for relevant work. Third, we map customisable criteria onto Goal Tokens to automate summarization and evaluation. We present "tangible prompting" as a framework for rethinking to our increasingly monotonous interactions with LLMs.

Authors Keywords

Tangible Interface, Human-AI Interface, Tangible Prompting

CSS Concepts

• Human-centered computing-> Human computer interaction (HCI) -> Interaction paradigms

INTRODUCTION

Conversational chat interfaces have rapidly become the dominant interface between humans and large language models (LLMs). While easy to understand, research on chat-based systems suggests they encourage passive consumption rather than active reflection. Users frequently cognitively disengage from generated outputs, skim content, overlook details, sometimes even forget what prompted them, and lose the thread of their own thinking [12].

This limitation becomes especially prominent in AI-assisted creative thinking. There are good reasons designers reach for a whiteboard rather than a word processor when a problem feels hard. Spatial tools like Miro demonstrate that non-linear, visual thinking environments unlock modes of reasoning that sequential text cannot — ideas coexist without hierarchy, connections emerge from proximity, and the whole composition remains visible at once. Current LLM interfaces enforce a cognitive queue fundamentally at odds with how creative thinking unfolds.

Especially in a time when artificial intelligence and automation make intelligence and productivity seemingly cheaper than ever, the skills that remain distinctly human – creative synthesis, critical judgment, lateral thinking – become ever more important to cultivate. Knowing which parts of your thinking to share with an AI, and which to develop yourself, is becoming one of the most important literacies of our time.

We present Tangible Co-Ideation, a spatial LLM interface paired with tangible objects that structure the prompting process through physical manipulation rather than typing. Drawing on embodied cognition, we aim to make brainstorming more playful and narrative-driven. The Tangible Co-Ideation begins with an idea. The user draws a quick sketch and describes their idea in a few words. Then the journey begins...

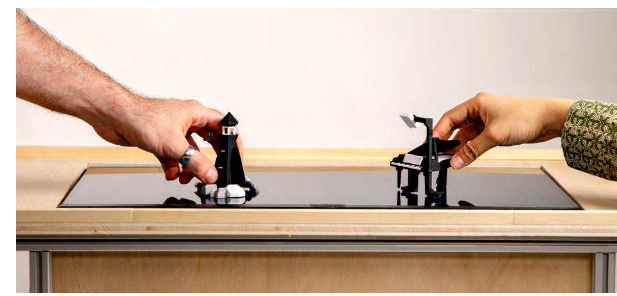

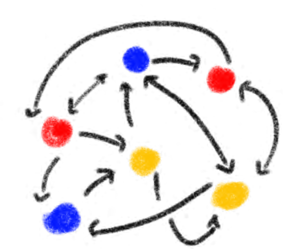

2. RETHINK

2.1 Cognitive Engagement through Embodiment

Recent studies found that LLM users showed severe deficits in critical thinking, coherent reasoning, and feelings of ownership [3,4,12], suggesting that passive, text-based AI interaction actively disengages users from their own thinking. This stands in contrast to what we know about physical interaction: tangible interfaces offer real benefits for learning and can simplify problem-solving [14]; natural learning engages multiple senses [20]; haptic and visual perception share abstract representations [5]; and physical objects are recalled better than pictures, and pictures better than words [2]. Although it remains debated whether these benefits stem from physicality itself or from the active engagement it affords [16, 19], what is consistent is that embodied interaction correlates with deeper cognitive processing and a way to keep the human meaningfully in the loop.

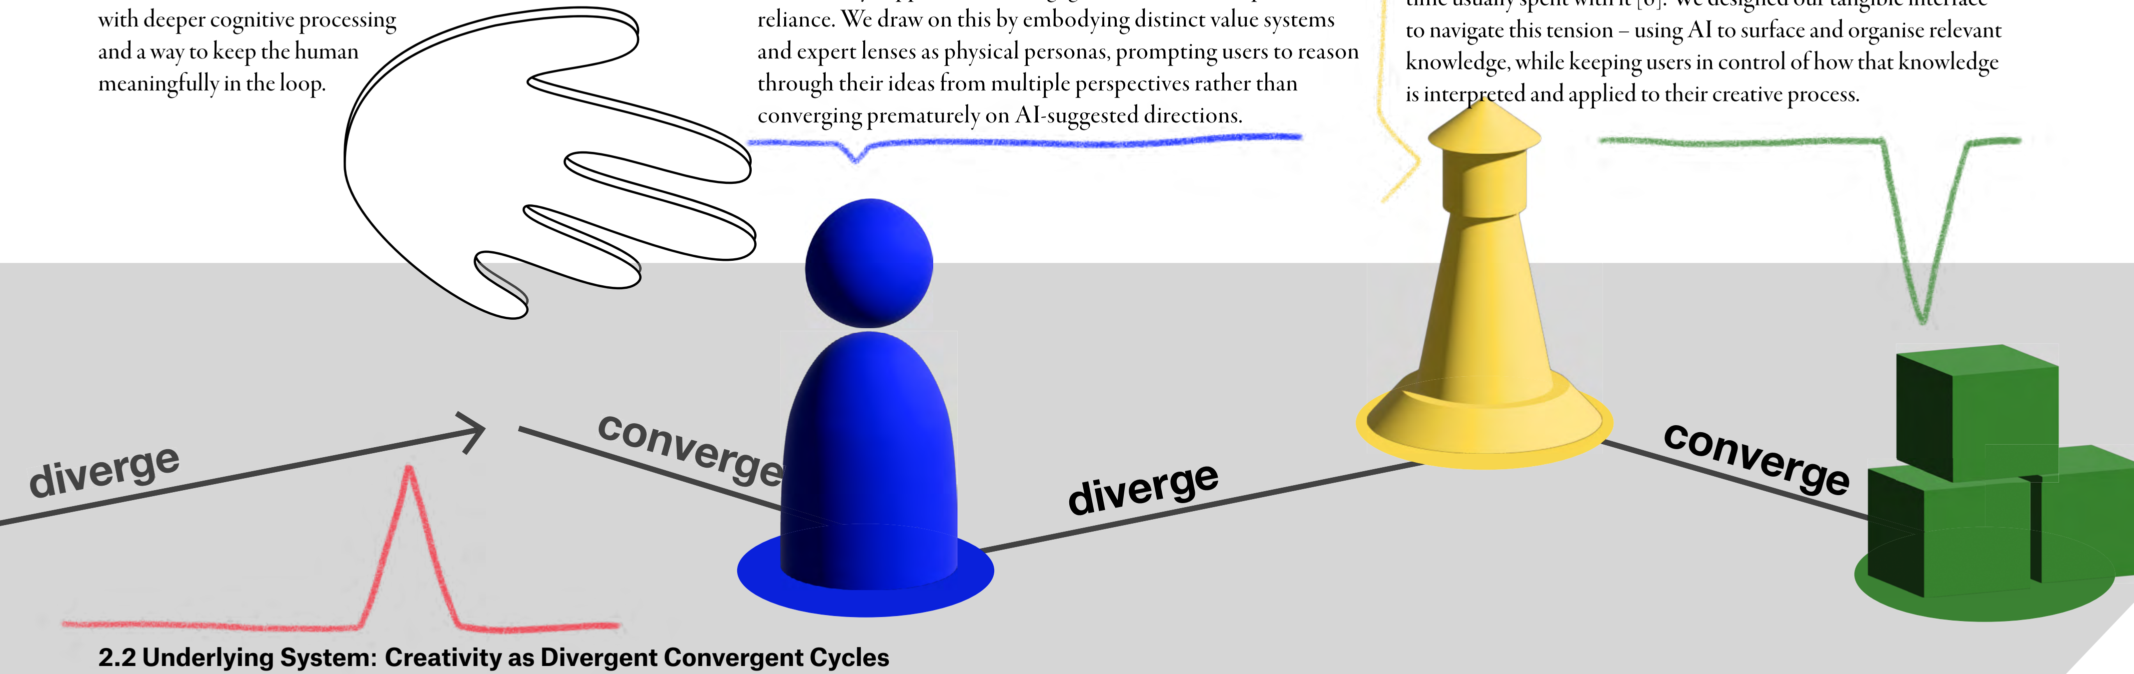

2.2 Underlying System: Creativity as Divergent Convergent Cycles

Creativity is widely understood as an iterative process alternating between divergent thinking (generating ideas) and convergent thinking (evaluating and refining them) [6,10]. HCI research on Creativity Support Tools confirms these phases are interdependent rather than sequential [7,8,19]. Supporting creativity, therefore, requires interaction structures that afford both expansion and evaluation within an evolving design space [23–25]. We designed our tangible interface around this: each physical object corresponds to a distinct phase of the creative cycle, allowing users to move fluidly between exploration and evaluation without the linear constraints of a chat interface.

2.3 AI Expert Personas

Generative AI has been shown to support both divergent and convergent thinking in creative processes, particularly for novices [2,17,21]. Tools like PersonaFlow demonstrate how LLM-simulated expert perspectives can enhance early-stage conceptual exploration [13], and AI has been shown to act as a sustained collaborative presence that shapes users’ conceptual strategies. However, AI-mediated ideation carries real risks: Shin et al. found no evidence that LLMs alone improve problem reframing and observed reduced agency among less-experienced users [15], while meaningful AI guidance requires careful calibration of trust. These findings underscore the need to design AI personas that actively support critical engagement rather than passive reliance. We draw on this by embodying distinct value systems and expert lenses as physical personas, prompting users to reason through their ideas from multiple perspectives rather than converging prematurely on AI-suggested directions.

2.4 AI for synthesizing large data

Particularly interesting for ideation is how LLMs excel at synthesizing large volumes of information through cross-referencing and associative knowledge generation across disparate sources [15]. Tools like Elicit demonstrate this for literature review. How can we leverage this skill in real time during ideation?

2.5 AI for summarizing and evaluation

Another area where LLMs shine is in automated meeting summary tools such as Otter.AI, or Fathom. They have been widely adopted to reduce cognitive overhead in professional workflows. Yet this convenience carries a risk: when AI handles synthesis, users risk losing meaningful contact with the underlying material and the time usually spent with it [6]. We designed our tangible interface to navigate this tension – using AI to surface and organise relevant knowledge, while keeping users in control of how that knowledge is interpreted and applied to their creative process.

3. REMAKE - AI PIPELINES FOR CREATIVE THINKING

AI Mentor

During Tangible Co-Ideation, the AI-Mentors serve as expert personas offering design critiques that represent different value systems, ranging from academic professors and artists to business-oriented entrepreneurs.

To develop AI-Mentors, we deployed a Retrieval-Augmented Generation (RAG) system, combined with supervised fine-tuning [13] and prompt engineering, to deliver precise, context-aware responses grounded in curated materials [1,18]. Relevant data, including academic papers, transcribed talks, sketches, and ChatGPT conversations, were collected, digitized, and indexed in a unified database [6]. The PDF of publications builds the foundation of the expert knowledge base. Recorded talks were transcribed to improve overall alignment and response accuracy. Sketches of human mentors served as a source of in-progress thinking. They were first processed using OCR and image analysis to extract textual and visual features. Then clustering algorithms and manual tagging were applied to organize sketches by topic and concept, and each visual element was linked to related documents to enable cross-modal retrieval [11,19]. Lastly, human mentors’ ChatGPT conversation logs were used to understand communication style. To ensure confidentiality, we built a custom application tool to

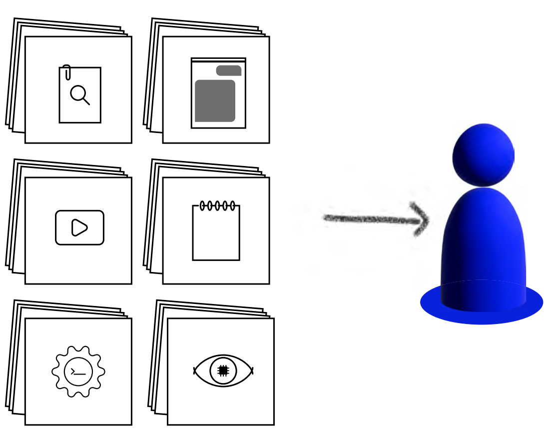

clean conversation logs. This tool removes irrelevant and personal exchanges under human supervision, resulting in a dataset for model training. Based on this curated data, AI-Mentor agents with defined expertise, background, and communication style were trained. In the last step, prompt engineering guided responses in the AI-Mentor’s voice and ensured appropriate referencing of documents and sketches. Prompts were iteratively tested and refined to optimize contextual grounding and response fidelity.

Lighthouse

The second feature is the Lighthouse (described in section 5.3), which helps users to search for relevant work during brainstorming. It uses a ChatGPT API call to analyse and identify topic clusters related to the user’s idea. The clusters are semi-structured as main, contextual, and technical. Then we use Semantic Scholar API, arXiv, and Crossref to find at least 5 publications per cluster. These more specific tools can better accommodate users’ imaginative ideas without leading to hallucinations.

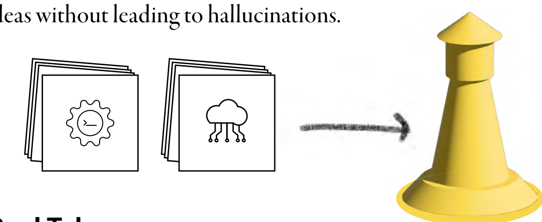

Goal Tokens

The third feature, Goal Tokens (described in section 5.4), uses LLMs as an automated summary and evaluation tool. To wrap up an ideation session, users place Goal Tokens that represent different criteria, e.g., originality, publishability, commerciality, etc., onto the field. Through prompt engineering, a summary of the brainstorming session, structured as advantages and disadvantages, as well as next steps, is presented to the user.

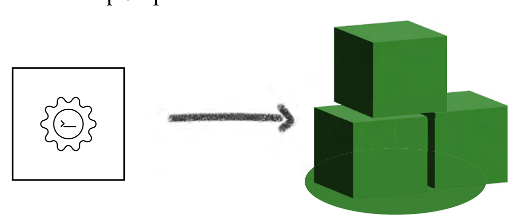

3.2 Capactive Sensing for Tangible Interaction

In our first prototype, we implemented an ArUco marker tracking system with an overhead camera. This approach enabled reliable identification of object identity, height, and 3D orientation. However, it required the top surface of each tangible to remain visible for tracking. Because our system emphasizes expressive, embodied manipulation, this constraint negatively impacts the interaction experience. In our second prototype, we explored a conductive filament–based approach. We 3D printed tangibles with three conductive contact points on their base, encoding object identity and orientation through the relative distances between these points. When placed on a capacitive touch screen, the system distinguishes tangibles based on their unique footprint geometry.This approach allows more natural interaction and works on most standard capacitive touch screens. A limitation is that most screens support only ten simultaneous touch points, restricting tracking to three tangibles at once. In practice, this was manageable, as our interaction scenarios rarely required more than three concurrent objects, and simple software strategies helped maintain object state. To preserve visual expressiveness, we carefully designed each tangible so that only necessary conductive regions remain exposed, while the main body uses non conductive colored materials. We generated and evaluated ten distinct footprint configurations with sufficient spatial separation to ensure reliable differentiation during interaction.

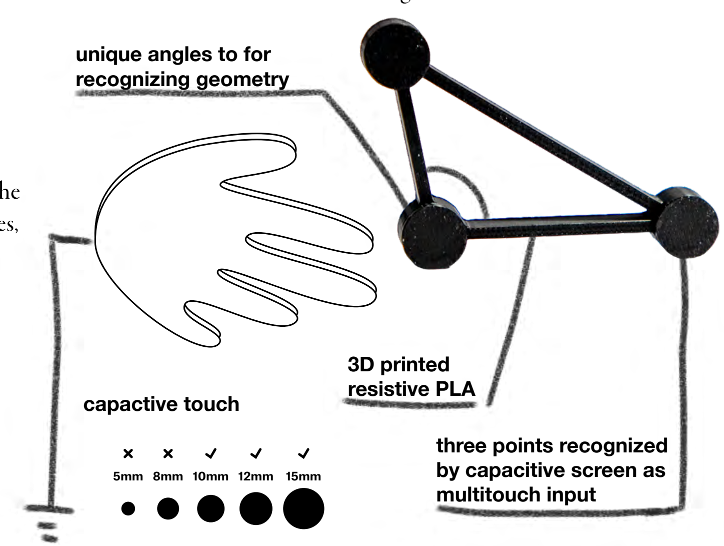

4. REMAKE - ITERATIONS

Software Design

The interface went through two design iterations. Each one addressed concrete problems identified during use. In the **first iteration**, we established the concept of "rooms" to structure tangible co-ideation. Ideas start with user input in rooms, where AI Mentors can join to give feedback. Rectangular rooms serve as containers for individual ideas. Each room could contain any combination of text, images, and freehand sketches, providing the context to prompt the mentors. Users could open multiple conversations per room and manually merge context from different threads to steer the generation.

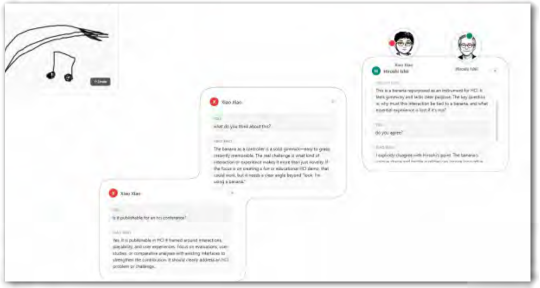

In the **second iteration**, we tried to embed an ideation sequence into the interaction, e.g., when the mentor was moved into a room, 3 bubble ideas were generated. Idea bubbles could be merged. When a mentor moved on an idea, it would emit a speech bubble with the mentor's opinion on the idea, and a proposal of a new direction. A lens object lets users visualize a live sketch of an idea.

A backpack mechanism allowed storing and carrying conversational threads across the canvas. One key limitation

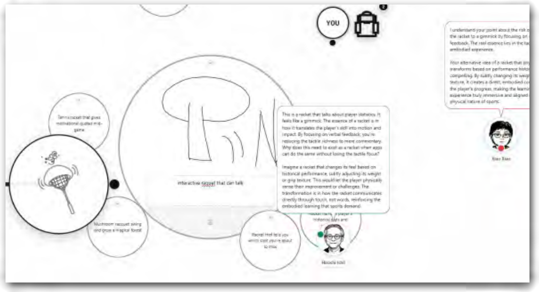

became clear through use: the many functions made it difficult to intuitively understand where to focus and how to use the system.

Hardware Design

The system includes two categories of physical objects with different design logics. The tokens (the Lighthouse and the Goal Tokens) are designed to draw from established conceptual metaphors [25, 26]: a lighthouse tower that rotates and beams light is already understood as something that searches or reveals, and a shelf that can tip is already understood as something that holds a state that can be inverted. These objects do not require instruction because users carry prior embodied schemas for what they do [27, 30].

The mentor figurines work differently, since they represent a person, or more accurately, a slice of a person or a bigger value system. We started with a humanoid-looking representation; however, we quickly moved to more symbolic representations. In the case of Hiroshi Ishii AI-Mentor, rather than being a digital twin of the real person, our RAG system creates an expert persona who represents an academic value system. After several rounds of form-finding, we chose the abacus. This tangible calculating device is often referenced by Prof. Ishii as the embodiment of his research visions. We wanted to be transparent that these AI-Mentors do not represent the entirety of a person, but rather a very specific aspect, in this case, Prof. Ishii's research side. Moreover, with more abstract figurines, we hope to encourage playful imagination and storytelling to identify problems rather than solve them during brainstorming.

Learnings

When we brought the software and hardware together, we quickly realized that when the software is developed without a tangible interface in mind and vice versa, there are many integration issues that break the experience.

- Moving UI elements does not fit well with tangible objects
- Tangible UI needs to be clearly connected with the screen, e.g., digital shadow
- The objects should be designed with legs that look like they are coming out of the screen rather than a flat surface, which covers a lot of screen real estate
- The objects need to be easily graspable and stable for optimal capacitive recognition

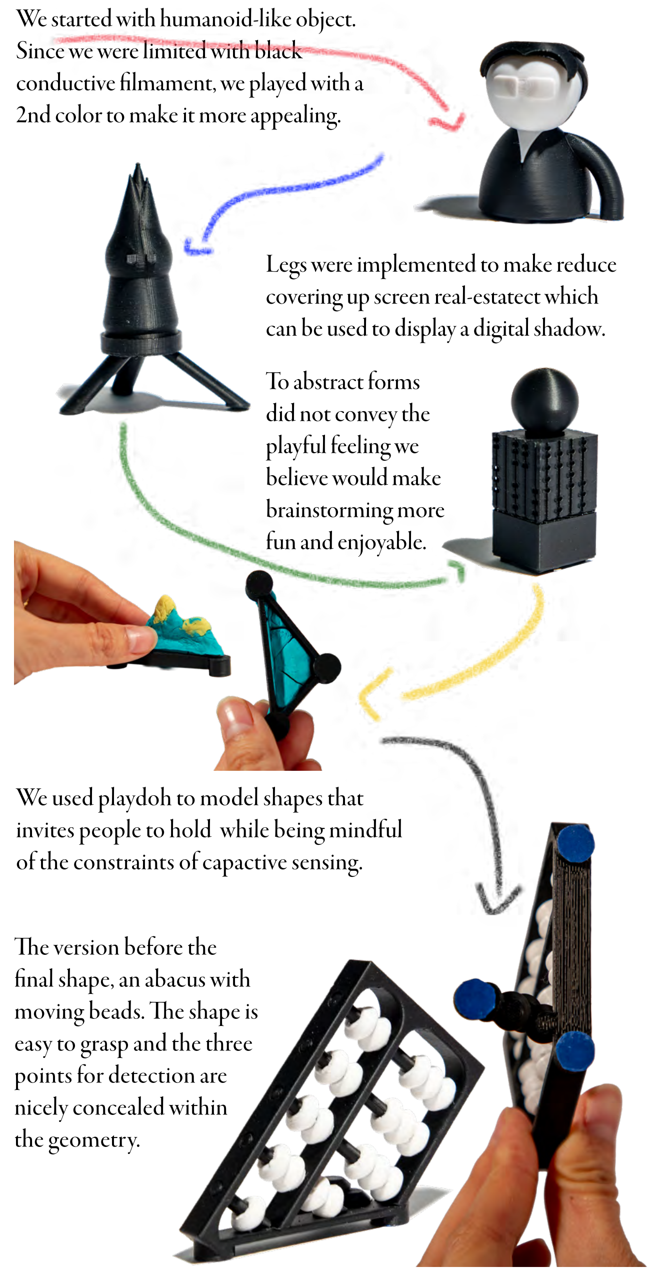

5. REDESIGN - TANGIBLE CO-IDEATION

5.1 Ideation mode

Tangible Co-ideation starts by drawing a circle on the screen. We call the circle a room, the starting point of an idea. The user then sketches an idea into the room and describes it in a few words.

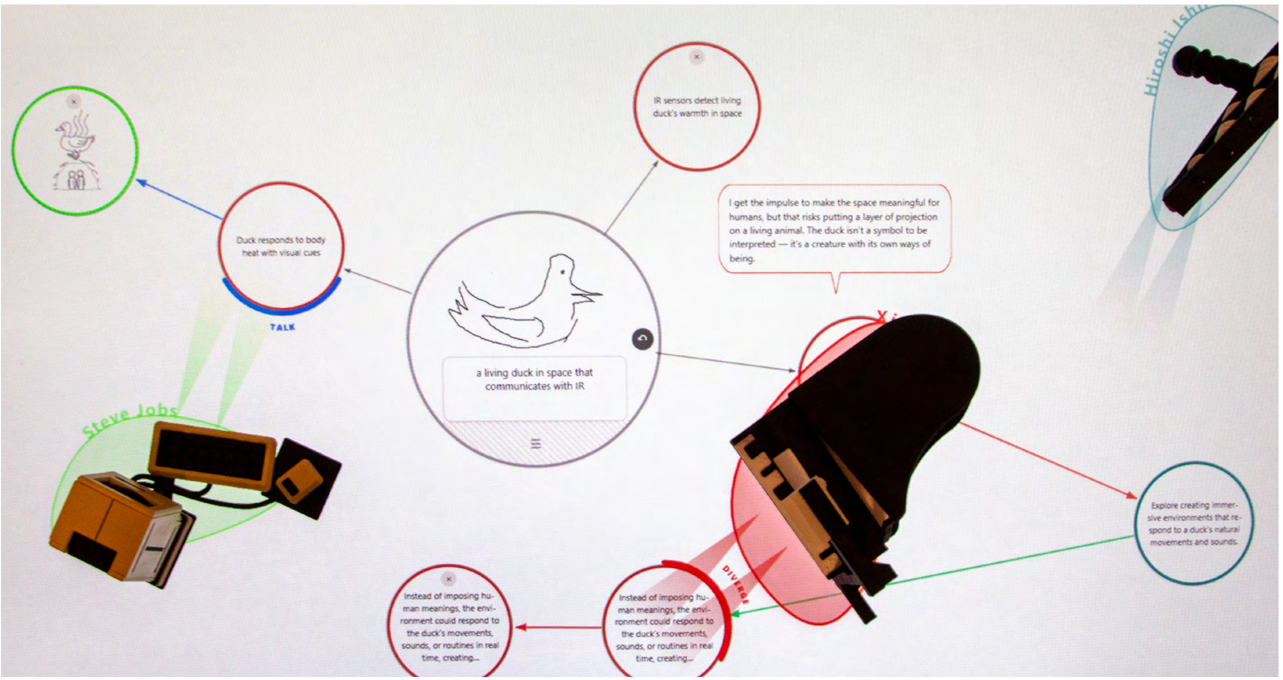

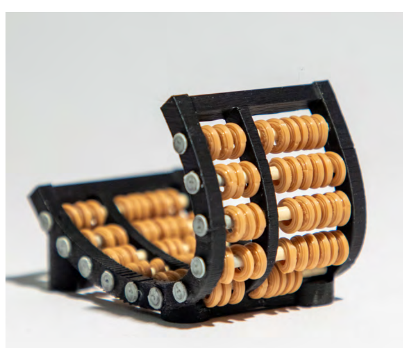

Next, an AI-Mentor figurine is physically placed in the room, which triggers a brainstorming sequence: the AI-Mentor reads the room's content and generates three ideas from its perspective. The ideas are shown as smaller circles, or bubbles, connected to the room. Each generated bubble is color-coded to indicate who said which comment and connects back to its source through a visible branch, so the chain of thought from the original room to any downstream idea stays traceable.

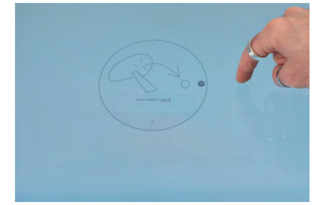

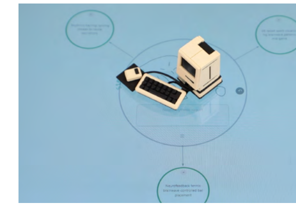

Next, the user picks one idea they are most intrigued by and guides the AI-Mentor figurine towards the bubble. Depending on the direction the mentor figurine approaches the bubble, the AI-Mentor will generate responses from different perspectives, converge, diverge, or respond to user-specific requests (talk mode). In converge mode, the mentor produces a response that builds toward the idea and links it to a new bubble derived from that convergent reasoning. In diverge mode, it challenges or reframes the idea and generates a bubble exploring an alternative direction. In talk mode, the user can direct the mentor with a specific question or request, including asking the mentor to sketch a response directly on the canvas. These perspectives challenge the user to make decisions and provide structure during brainstorming. The mentor's digital shadow and its animated digital gaze serve as visual indicators.

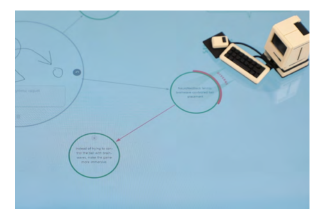

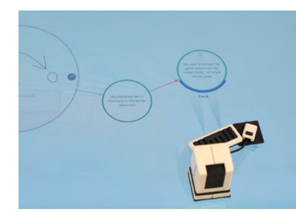

Multiple mentors can be active on the canvas at the same time and can each interact with the same ideas independently, simulating a design critique experience with experts from different perspectives. Furthermore, bubbles can be merged or used to instantiate new rooms to kick off a parallel or next iteration of brainstorming.

5.2 Debate Mode

Arena mode is triggered by a physical gesture: two mentor figurines placed facing each other, and the debate on the nearest idea bubble begins. The two mentors take turns generating responses, each arguing from the perspective of their trained persona. With two mentors, their spatial orientation continues to serve as tangible prompting language throughout the session:

- when both figurines face each other, the current speaker is prompted to agree and build on the previous point; when both turn away, the stance flips to disagreement.
- Asymmetric orientations produce one mentor trying to convince while the other refuses to engage. In addition, there are four physical gestures in debate mode.
- Moving a figurine back and forth horizontally, similar to shaking the head, triggers a sceptical gesture: the next response opens with doubt, challenges a specific claim, and refuses to accept the previous point.
- Tracing a circular motion with a figurine triggers an exciting gesture: the next turn becomes short, punchy, and increases enthusiasm for a thread of thought.
- Picking up a figurine and tapping it back down twice triggers an interruption: the current speaker is cut off midstream, and the tapping mentor immediately takes the floor.
- Physical contact between two figurines triggers a confrontation gesture: the next response opens with a blunt contradiction and escalates the disagreement directly.

These embodied gestures extend the argument that physical manipulation of objects can elicit a different level of cognitive engagement than verbal input. As with children playing with figurines, the tone, urgency, and stance become parameters that construct prompts.

The debate ends in one of two ways. Moving one of the figurines into the conclusion zone at the bottom-right corner of the screen and holding it there for three seconds closes the session. The mode can also auto-conclude if one mentor wins an argument over the other. On exit, a new room appears with a sketch and description, containing an AI-generated synthesis of the debate. Three idea bubbles are placed around it to kick off a new round of ideation.

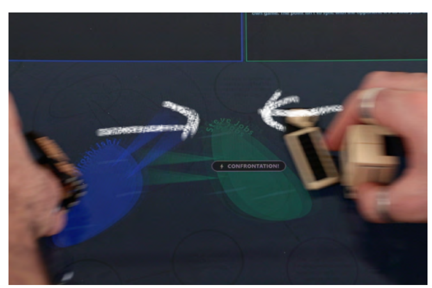

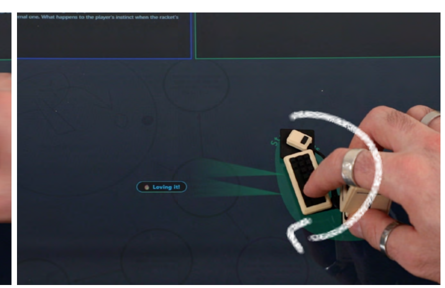

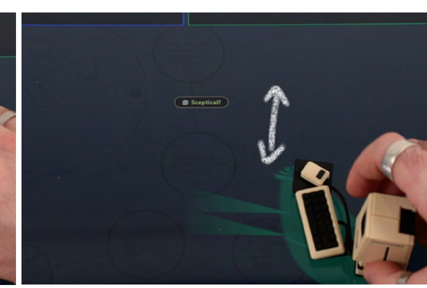

Two physical objects complement the mentor figurines.

5.3 Lighthouse Mode

A lighthouse guides ships through the dark by pointing in a direction. During brainstorming, the user is in a comparable position: generating ideas without knowing what already exists around them in the research landscape. When the lighthouse is placed on the surface, black smoke blurs the entire canvas. Overlayed on the ideas, four clusters appear. A 90-degree beam sweeps from the tower to the edge of the screen, cutting through the darkness and illuminating whichever cluster it lands on in yellow. Rotating the physical tower rotates the beam, letting the user sweep through the canvas and choose which cluster to surface.

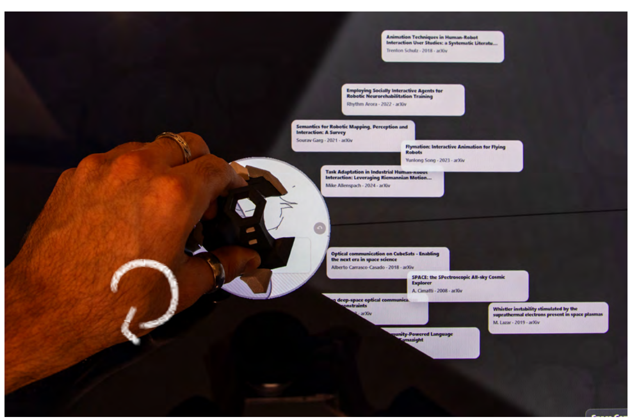

5.4 Goal Tokens Mode

Two Goal Tokens, represented through shelves, sit on the surface alongside the mentor figurines and lighthouse: a bookshelf for publishability and a supermarket shelf for commercialization. Placing either on an idea bubble immediately starts an evaluation sequence. For the publishability shelf, reports appear as notepads. For the supermarket shelf, they appear as receipts. The format inside each is fixed: four sections of three short bullets each. For the bookshelf, those sections are Strengths, Research questions, Hypotheses, and To develop. For the supermarket shelf: Market size, Target audience, Revenue potential, and Go-to-market. Tipping the shelf over regenerates all reports under a negative framing: the publishability sections become Weaknesses, Missing elements, Logical flaws, and Not publishable because. The commercialization sections become Market barriers, Target issues, Revenue risks, and Fatal flaws. The Goal Tokens summarize a brainstorming session by presenting its strengths and weaknesses through the lens of set goals. It helps users to converge and start the next round of ideation with continuity.

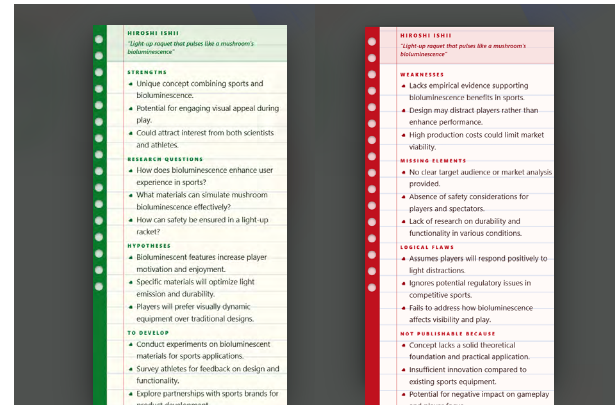

CONCLUSION

In today's educational landscape, AI is an inseparable component. Students rapidly learn how to use these tools to their advantage. The risk is not that they will use AI poorly. It is that they will use it so fluently that many don't realize their own thinking might quietly disappear. Knowing how to use AI to think more clearly and which tasks should be done yourself is becoming one of the most important literacies of our time.

Creative synthesis, critical judgment, and the capacity to hold an unresolved problem long enough to understand it are not skills fostered by current Human-AI interfaces. They require time, effort, and friction. Tangible Co-Ideation is built on these premises. Every interaction requires a physical act: placing a mentor, tapping figurines, rotating a lighthouse, tipping a shelf, or moving two mentors toward each other. Spatial relationships are leveraged to prompt in tangible ways. The angle a figurine faces shapes what gets generated. The direction the lighthouse points determines what literature surfaces; the orientation of two mentors mid-debate edits the tone of the argument in real time. Shneiderman argued at CHI 1998 that the pursuit of transformational technologies must remain oriented toward human flourishing rather than mere automation. That question has only become more urgent in the decades since. "Our greatest resource is human creativity. Appropriate technology can empower more people to experience the thrill of creative work and enable each of us to make a difference" [16].

Future: seamless implementations across hardware platforms with capactive sensing capabilities

REFERENCES

[1] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. https://doi.org/10.48550/arXiv.2310.11511

[2] Benett Axtell, Raheleh Saryazdi, and Cosmin Munteanu. 2022. Design is Worth a Thousand Words: The Effect of Digital Interaction Design on Picture-Prompted Reminiscence. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22), 1–12. https://doi.org/10.1145/3491102.3517692

[3] Valdemar Danry, Pat Pataranutaporn, Yaoli Mao, and Pattie Maes. 2020. Wearable Reasoner: Towards Enhanced Human Rationality Through A Wearable Device With An Explainable AI Assistant. In Proceedings of the Augmented Humans International Conference (AHs '20), 1–12. https://doi.org/10.1145/3384657.3384799

[4] Valdemar Danry, Pat Pataranutaporn, Yaoli Mao, and Pattie Maes. 2023. Don't Just Tell Me, Ask Me: AI Systems that Intelligently Frame Explanations as Questions Improve Human Logical Discernment Accuracy over Causal AI explanations. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23), 1–13. https://doi.org/10.1145/3544548.3580672

[5] Randolph D. Easton, Anthony J. Greene, and Kavitha Srinivas. 1997. Transfer between vision and haptics: Memory for 2-D patterns and 3-D objects. Psychonomic Bulletin & Review 4, 3: 403–410. https://doi.org/10.3758/BF03210801

[6] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2025. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. https://doi.org/10.48550/arXiv.2404.16130

[7] Ronald A. Finke, Steven M. Smith, and Thomas B. Ward. 1992. Creative Cognition: Theory, Research, and Applications. MIT Press, Cambridge, MA, USA.

[8] 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 2019 CHI Conference on Human Factors in Computing Systems (CHI '19), 1–18. https://doi.org/10.1145/3290605.3300619

[9] Jonas Frich, Michael Mose Biskjaer, and Peter Dalsgaard. 2018. Twenty Years of Creativity Research in Human-Computer Interaction: Current State and Future Directions. In Proceedings of the 2018 Designing Interactive Systems Conference, 1235–1257. https://doi.org/10.1145/3196709.3196732

[10] J. P. Guilford. 1950. Creativity. American Psychologist 5, 9: 444–454. https://doi.org/10.1037/h0063487

[11] Geewook Kim, Teakgyu Hong, Moonbin Yim, JeongYeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoo Yun, Dongyoon Han, and Seunghyun Park. 2022. OCR-Free Document Understanding Transformer. In Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVIII, 498–517. https://doi.org/10.1007/978-3-031-19815-1_29

[12] Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. 2025. Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. https://doi.org/10.48550/arXiv.2506.08872

[13] 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), 506–534. https://doi.org/10.1145/3715336.3735789

[14] Claire O'Malley and Danae Stanton Fraser. 2004. Literature Review in Learning with Tangible Technologies. Retrieved February 23, 2026 from https://telearn.hal.science/hal-00190328

[15] 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, 1–25. https://doi.org/10.1145/3706598.3713273

[16] Ben Shneiderman. 1998. Codex, memex, genex: the pursuit of transformational technologies. In CHI 98 Conference Summary on Human Factors in Computing Systems (CHI '98), 98–99. https://doi.org/10.1145/286498.286548

[17] Ben Shneiderman. 2007. Creativity support tools: accelerating discovery and innovation. Communications of the ACM 50, 12: 20–32. https://doi.org/10.1145/1323688.1323689

[18] Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. 2024. Corrective Retrieval Augmented Generation. https://doi.org/10.48550/arXiv.2401.15884

[19] Shi Yu, Chaoyue Tang, Bokai Xu, Junbo Cui, Junhao Ran, Yukun Yan, Zhenghao Liu, Shuo Wang, Xu Han, Zhiyuan Liu, and Maosong Sun. 2025. VisRAG: Vision-based Retrieval-augmented Generation on Multi-modality Documents. https://doi.org/10.48550/arXiv.2410.10594

[20] Oren Zuckerman, Saeed Arida, and Mitchel Resnick. 2005. Extending tangible interfaces for education: digital montessori-inspired manipulatives. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05), 859–868. https://doi.org/10.1145/1054972.1055093

[21] A morphological analysis of the design space of input devices | ACM Transactions on Information Systems. Retrieved February 27, 2026 from https://dl.acm.org/doi/10.1145/123078.128726

[22] 22. A Design Space for Gaze Interaction on Head-mounted Displays | Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. Retrieved February 27, 2026 from https://dl.acm.org/doi/10.1145/3290605.3300855

[23] Filtering and Informing the Design Space: Towards Design-Space Thinking: ACM Transactions on Computer-Human Interaction: Vol 28, No 1. Retrieved February 27, 2026 from https://dl.acm.org/doi/10.1145/3434462 https://dl.acm.org/doi/10.1145/3434462

[24] Ben Shneiderman. 1998. Codex, memex, genex: the pursuit of transformational technologies. In CHI 98 Conference Summary on Human Factors in Computing Systems (CHI '98). Association for Computing Machinery, New York, NY, USA, 98–99. https://doi.org/10.1145/286498.286548

[25] Lakoff, G. and Johnson, M. (1980). Metaphors We Live By. University of Chicago Press.

[26] Johnson, M. (1987). The body in the mind: The bodily basis of meaning, imagination, and reason. University of Chicago Press.

[27] Gibson, J.J. (1979). The Ecological Approach to Visual Perception. Houghton Mifflin.

[28] Hutchins, E. (1995). Cognition in the Wild. MIT Press.

[29] Slamecka, N. J., & Graf, P. (1978). The generation effect: Delineation of a phenomenon. Journal of Experimental Psychology: Human Learning and Memory, 4(6), 592–604. https://doi.org/10.1037/0278-7393.4.6.592

[30] Donald A. Norman. 1999. Affordance, conventions, and design. interactions 6, 3 (May/June 1999), 38–43. https://doi.org/10.1145/301153.301168

[31] Michelene T. H. Chi & Ruth Wylie (2014) The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes, Educational Psychologist, 49:4, 219-243, doi: 10.1080/00461520.2014.965823