education sciences

MDPI

Article

Intermediate Representations in Human–AI Creative Collaboration: A Framework for Distributed Creative Cognition

Bruce Donald Campbell [1,*]

[1] Rhode Island School of Design; bcampbel01@risd.edu
* Correspondence: bcampbel01@risd.edu

Abstract

Generative artificial intelligence has dramatically reduced the effort required to produce visual artifacts, intensifying questions about which human practices remain educationally valuable when computational systems can generate images almost instantly. Much discussion of human-AI creative collaboration begins with text-based prompting. This conceptual paper argues that collaboration often begins earlier, while intentions are still being formed through observation, sketching, reflection, and other intermediate representations. The framework developed here was motivated by reflective teaching practice and observation across repeated offerings of an elective course in art and design education. Classroom experiences are used as illustrative context rather than as evidence of causal effects. Drawing on embodied cognition, distributed cognition, distributed intelligence, and boundary-object theory, the paper conceptualizes sketches as evolving cognitive interfaces that externalize partially formed ideas, preserve productive ambiguity, and provide reference points against which AI interpretations can be accepted, rejected, or revised. Two complementary models describe a sketch-mediated creative cycle and the broader distributed cognitive system in which that cycle operates. The framework suggests that the educational value of sketching may increasingly lie not in artifact production but in supporting intention formation, reflective judgment, human agency, and productive friction before and during AI interaction. The paper concludes with propositions for future empirical research on sketch-first and other representation-mediated forms of human-AI collaboration.

Keywords: embodied cognition; distributed cognition; distributed intelligence; intermediate representations; human-AI collaboration; art and design education; creativity; sketching

1. Introduction

The emergence of generative artificial intelligence has transformed creative practice. Tasks that once required substantial technical training can now be accomplished by way of natural-language prompts in seconds. This capability has generated both excitement and uncertainty within educational institutions. If students can produce compelling visual artifacts without drawing, modeling, or painting, what educational purpose remains for these foundational practices? More fundamentally, what forms of human learning become increasingly valuable as intelligent systems assume greater responsibility for producing artifacts?

Much educational discussion surrounding generative AI has focused on academic integrity, assessment, prompt engineering, and responsible use. These concerns are important, but they often imply that human-AI collaboration begins with a prompt. Within that view, the educational challenge becomes helping students formulate effective instructions for increasingly capable computational systems. This paper explores a different possibility: meaningful human-AI collaboration can begin before prompting, during the formation of ideas themselves.

Creative work rarely begins with fully articulated intentions. Designers, artists, scientists, engineers, and architects often work through incomplete observations, tentative relationships, and provisional concepts. Sketches, diagrams, notes, models, and other external representations do more than record thinking; they can participate in its development (Goel, 1995; Kirsh, 2010). Sketches are especially useful because their incompleteness can preserve ambiguity while ideas remain uncertain (Goldschmidt, 1991).

This distinction becomes important in the presence of generative AI. Image-generation systems can respond to articulated instructions and visual inputs, but the human processes through which intentions emerge still precede and surround those interactions. As AI accelerates artifact production, educational attention may therefore need to shift toward observing carefully, externalizing incomplete ideas, reflecting on emerging representations, and making judgments about computational suggestions.

The framework proposed in this paper grew from a reflective teaching practice in an elective course on artificial intelligence in art and design education (Campbell, 2025). Across six offerings of the course, sketch-first activities repeatedly raised questions about how students formed intentions, evaluated AI outputs, revised ideas, and described ownership of their creative processes so others could understand their collaboration with AI. These classroom experiences are not presented as a controlled comparison or as evidence that sketching causes particular learning outcomes. They serve instead as the motivating context for a conceptual question: How might intermediate representations shape human learning and collaboration in the presence of increasingly capable artificial intelligence?

The contribution of this paper is a framework for understanding sketches as intermediate representations within distributed human-AI creative cognition. The argument is not that sketching necessarily improves AI-generated imagery. Rather, sketching can give partially formed human intention an external form before AI contributes its own interpretation. This creates a representational space in which perception, reflection, ambiguity, computational generation, critique, and revision can interact. Although the discussion begins with drawing, the framework is intended to apply more broadly to other representations through which humans think before and while collaborating with intelligent systems.

2. Conceptual Foundations

The framework brings together four complementary perspectives. Each treats cognition as more than the production of an internal idea followed by its external expression. Together they help explain why an apparently traditional practice such as sketching may acquire renewed importance in AI-assisted creative work.

2.1 Embodied cognition and sketching as thought

Sketching joins perception and action. A person observes, moves a hand, sees the resulting mark, and responds to what has become visible. Relationships may be discovered during drawing rather than fully specified beforehand (Suwa & Tversky, 2009). This is consistent with embodied approaches in which cognition emerges through interaction among mind, body, and environment (Varela et al., 1992). From this perspective, drawing

is not simply the transcription of a completed internal idea. The activity can participate in forming the intention that will later be communicated to an AI system.

2.2 Distributed cognition

Once a sketch exists, part of the developing idea persists outside the individual. It can be inspected, compared, annotated, discussed, and revisited. Distributed cognition describes cognitive processes that extend across people, artifacts, and environments (Hutchins, 1996). Human-AI creative work can therefore be understood as a cognitive system comprising the learner, intermediate representations, computational outputs, tools, peers, instructors, and the surrounding task environment. AI becomes a participant in this system rather than a replacement for the human processes that precede it.

2.3 Distributed intelligence and intermediate representations

Distributed intelligence emphasizes how artifacts and representations reorganize intellectual activity (Pea, 1993). Sketches occupy a productive middle ground: they are more persistent than an unexternalized thought but less fixed than a finished artifact. Their incompleteness allows ideas to remain revisable. Similar functions are served by mathematical notation, engineering drawings, scientific visualizations, conceptual maps, musical notation, and computational models. The central claim of this paper is therefore broader than drawing: intermediate representations can become cognitive interfaces through which biological and artificial intelligences coordinate evolving work.

2.4 Boundary objects

Sketches can also function as boundary objects (Star & Griesemer, 1989). A single drawing can support reflection by its creator while also being interpreted differently by classmates, instructors, and AI systems. It remains stable enough to coordinate attention yet ambiguous enough to permit alternative interpretations. In human-AI collaboration, discrepancies among intended, sketched, and generated forms need not be treated simply as failures. They can expose assumptions, prompt clarification, or open a new conceptual direction. The value lies partly in making those differences available for judgment.

These perspectives converge on a common idea: an intermediate representation does not merely carry information from a human mind to an AI system. It can reorganize human cognitive activity while also altering the inputs and computational processes through which the AI system generates a response, providing a persistent object around which subsequent interaction and interpretation are organized.

2.5 Contemporary human–AI co-creativity with intermediate representation

Recent research on generative AI has increasingly moved beyond treating artificial intelligence as an automated production tool toward examining human–AI co-creativity as an interactive process. Rezwana and Maher's Co-Creative Framework for Interaction Design (COFI), for example, emphasizes interaction dynamics such as turn-taking, contribution type, and communication between human and computational partners. Their analysis shifts attention from the capabilities of the AI alone toward the structure of the interaction through which co-creation occurs (Rezwana & Maher, 2023). This perspective is compatible with distributed accounts of cognition, and it still raises an important question for creative education: what cognitive and representational activity occurs before the AI becomes a participant in that interaction?

This question becomes particularly important when human agency is considered. Generative systems can extend ideation and provide alternatives that might not otherwise occur to a designer, but AI assistance can also change which parts of a task remain under

active human control. One large randomized study of AI-supported peer feedback found that students tended to rely on AI assistance rather than learn sufficiently from it to maintain the same performance when that assistance was removed (Darvishi et al., 2024). Although their study looked at peer feedback rather than creative design, its implications for agency are relevant here: effective AI assistance and development of human capability are not necessarily equivalent educational outcomes. This distinction suggests that human–AI collaboration should be evaluated not only by what the combined system produces, but also by the cognitive activity retained by the human participant.

One prominent response to the challenge of effective human–AI interaction has been increased attention to AI literacy and prompt engineering. Knoth et al. connect prompt-engineering strategies explicitly to AI literacy and not merely as a technical shortcut (Knoth et al., 2024). Oppenlaender et al. characterize prompt engineering for text-to-image generation as a creative skill that requires learning and practice. Their studies show that prompting itself is neither trivial nor cognitively empty, and that users must learn how to translate intentions into forms that generative systems can interpret effectively (Oppenlaender et al., 2023).

The framework in figures 1 and 2 does not challenge the importance of prompt competence. Instead, it directs attention to an earlier stage of the process. A prompt is already an external representation of an intention. By the time a learner formulates a prompt, some decisions about what matters, what should be specified, and what can remain ambiguous have already been made. Creative intentions, however, do not necessarily originate in linguistic form. They may initially be perceptual, spatial, gestural, relational, or only partially articulated. From this perspective, prompt engineering addresses the problem of communicating an intention to an AI system, while sketch-mediated interaction also asks how that intention becomes sufficiently formed for communication in the first place.

Recent research on generative AI in design provides evidence for the value of keeping this distinction visible. Paananen et al., studying text-to-image generation during architectural ideation, found that generative tools could support serendipitous discovery and an imaginative mindset, while also identifying the importance of design constraints and designers' ability to work with the limitations of the systems (Paananen et al., 2024). Such findings illustrate a productive tension in AI-supported design: generative systems can expand the space of possibilities, while designers must retain sufficient intention and judgment to determine which computational possibilities are relevant to the developing design.

The increasing availability of multimodal generative systems makes the role of sketches especially significant. Edwards et al., for example, developed Sketch2Prototype, a framework that begins with a hand-drawn sketch and moves through text, image, and three-dimensional representations (Edwards et al., 2024). Their results also highlight the value of intermediate modalities within an iterative design process. More recently, Baudoux et al. developed a multimodal system specifically supporting text-based and sketch-based human–AI conversations during conceptual design (Baudoux et al., 2025). Their work treats multimodal interaction as a means of supporting more agile and naturalistic exchanges between designers and generative systems. Together, these studies demonstrate a growing movement beyond text-only prompting toward interaction in which representations can move among sketches, language, generated images, and other modalities.

However, the use of a sketch as an input modality and its use as a cognitive intermediate representation are not identical. In multimodal AI research, a sketch can be evaluated according to what it communicates to a generative system or how effectively it constrains subsequent output. The framework proposed here emphasizes a complementary function. Before the sketch communicates anything to the AI, producing it may already have changed what the human perceives, intends, notices, and is prepared to evaluate.

The sketch therefore participates in two relationships: one between the learner and the developing idea, and another between the learner and the generative system.

This distinction also changes how the speed and timing of generative AI can be interpreted educationally. Rapid generation is valuable for exploration because it allows designers to encounter alternatives quickly. Yet speed can also remove activities that previously occurred while an idea was being externalized. Recent empirical work suggests that the timing of AI participation may itself influence learning processes. Yang et al., for example, incorporated strategically delayed generative AI into a collaborative design-thinking environment for preservice teachers and found the AI most effective when introduced after human-only design efforts (Yang et al., 2026). Their intervention differs substantially from the sketch-mediated practices considered here, but the result encourages a broader question raised by this framework: what cognitive work should occur before AI enters a creative process? Drawing requires a learner to spend time observing, selecting, positioning, simplifying, and revising. From the perspective developed here, this delay need not be treated simply as inefficiency. It may constitute productive friction through which an emerging intention becomes available for reflection before computational alternatives are introduced.

Recent research suggests that creative agency in human–AI collaboration is not established once but must be sustained throughout the creative process. Guo et al. found that independent thinking before or after AI collaboration could strengthen human agency, but that deeper integration of human ideas with AI-generated content could better sustain agency during collaborative creative problem solving (Guo et al., 2025). Complementing this finding, Rafner et al. describe creative agency as dynamic and fluctuating across human–AI co-creation, involving creative self-efficacy, control, autonomy, and ownership of the resulting work. Their participants used strategies such as progressive refinement and selective appropriation to maintain or recover agency as AI contributed unexpected material (Rafner et al., 2025). Together, these findings suggest that preserving agency involves more than establishing human intention before AI interaction; it requires continued opportunities to evaluate, integrate, reject, and reshape computational contributions. A revisable intermediate representation such as a sketch may provide one mechanism for aiding this continuing negotiation of agency and ownership.

The argument is not that sketching should replace prompting, nor that manual representation is inherently more creative than AI-assisted production. Rather, the timing and form of AI participation become variables in the organization of creative cognition. The sketch-first framework developed in this paper connects these strands by asking what cognitive work an intermediate representation performs for the human before, as well as during, its use in communication with AI.

3. From Classroom Practice to a Conceptual Framework

A framework was motivated by six offerings (2024-2026) of Understanding and Exploring Artificial Intelligence (UEAI), an elective course in an art and design curriculum. Sixty-five students from varied creative disciplines encountered generative AI alongside observational and conceptual sketching. Students were encouraged to treat AI as a participatory collaborator so as to find useful creative processes through that lens. Instruction emphasized process over product and repeatedly used cycles of observation, sketching, reflection, AI interaction, critique, and revision.

Several recurring classroom experiences prompted the present conceptualization. When a sketch existed before prompting, students had an external representation against which an AI interpretation could be judged. AI outputs could then function as additional design material rather than automatically becoming endpoints. Students often moved back and forth among drawing, prompting, discussion, and revision. Sketches also gave

peers and the instructor a common object for discussing intentions that were difficult to express verbally. In some cases, an unexpected AI interpretation led to further drawing or to a change in the student's conception of the work.

These experiences did not occur uniformly. A student could abandon an initial sketch, prefer an AI interpretation, or discover that an early representation constrained rather than expanded exploration. Such cases are important because the framework does not assume that sketching always improves collaboration. Instead, they suggest that the role of an intermediate representation depends on how it is used: as a provisional expression of intention, a constraint, a reference for critique, or an object that itself becomes open to revision.

Figure 1 summarizes the recurring process that motivated the framework. It is not a prescribed instructional sequence and should not be read as a well-tested process model. Its purpose is to make visible a key departure from prompt-centered accounts of human-AI collaboration: substantial cognitive activity may occur before the first prompt, and AI output can return the learner to representation and reflection rather than terminating the process.

Figure 1. Sketch-mediated creative cycle observed in instructional practice. The figure is an interpretive synthesis used to motivate the conceptual framework rather than a prescribed or experimentally validated sequence.

4. Sketch-Mediated Human-AI Collaboration

Figure 2 expands the process into a distributed cognitive system. The sketch occupies an intermediate position among perception, reflection, communication, and generation. It is neither merely an input format nor necessarily the origin of the idea. Instead, it is a

temporary stabilization of an evolving intention that can be encountered by the learner and by other participants, including AI.

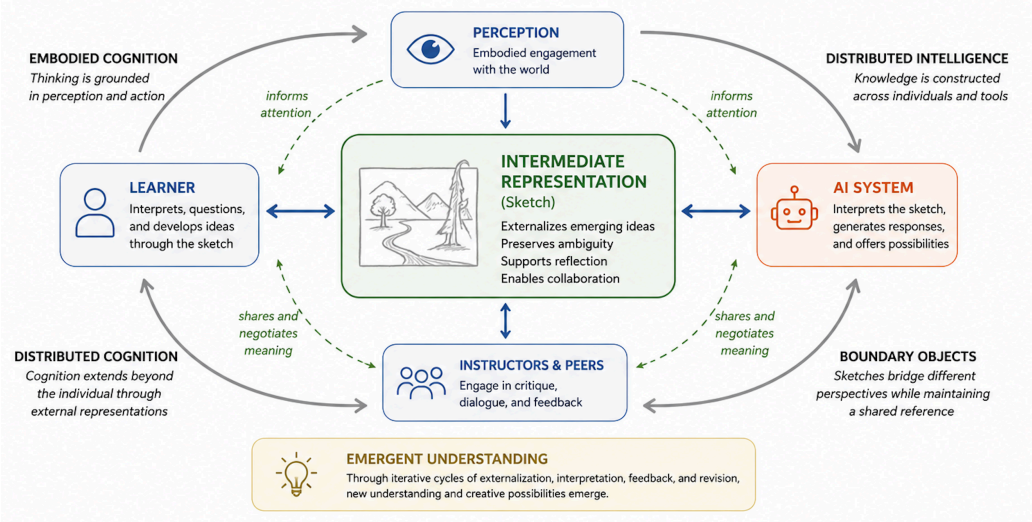

Figure 2. Conceptual framework for sketch-mediated human-AI collaboration. Sketches function as intermediate representations coordinating perception, reflection, communication, and generation across learners, peers, instructors, and AI systems. The framework draws on embodied cognition (Varela et al., 1992), distributed cognition (Hutchins, 1996), distributed intelligence (Pea, 1993), and boundary objects (Star & Griesemer, 1989).

4.1 Intention formation before prompting

Prompt-centered interaction encourages attention to how an already formed intention should be expressed in language. Sketch-mediated interaction adds an earlier problem: how does the intention become sufficiently formed to communicate at all? Observational and conceptual drawing can make relationships visible before they are easily verbalized. The sketch therefore provides AI with more than visual information; it is the residue of prior human perception and decision-making.

4.2 Iteration as a representational negotiation

When AI responds to a sketch or to a prompt derived from one, the generated artifact introduces another representation into the system. The human can compare that representation with the sketch, the underlying intention, and newly emerging possibilities. Revision may occur in the prompt, the sketch, or the conception itself. Human-AI collaboration is therefore better described as representational negotiation than as a one-way transfer from prompt to output.

4.3 Ambiguity as a resource

A sketch can deliberately leave relationships unresolved. This ambiguity distinguishes many intermediate representations from specifications. AI may fill those gaps in ways that expose an assumption or propose an alternative the human had not considered. Productive collaboration does not require eliminating ambiguity; it requires the human to recognize which aspects are intentionally open and to judge what the AI contributes there. The same ambiguity can also create risk when computational defaults displace an intention that was never sufficiently externalized to be defended or revised consciously.

4.4 Human agency as continued judgment

Within this framework, agency does not require that the human determine every feature of the final artifact. It resides in continued responsibility for the developing vision. An intermediate representation can assist that responsibility by providing a human-originated reference against which computational suggestions are accepted, rejected, or transformed. This is a proposed mechanism, not a demonstrated effect. Whether sketch-first practices measurably strengthen perceived or observable creative agency remains an empirical question.

4.5 Productive friction

Generative AI makes production fast. Sketching makes some parts of the process deliberately slower. That friction may be educationally useful when it creates time for observation, uncertainty, comparison, and reflection. Efficiency is therefore an incomplete criterion for evaluating AI-supported creative learning. A workflow that requires more time may promote richer intention formation or clearer provenance of decisions, although those possibilities require direct study.

Increasingly capable multimodal AI systems may eventually reduce the instrumental need to sketch in order to communicate visual intentions. The present framework suggests, however, that communicative efficiency and cognitive value are different questions: a representation may remain educationally useful because of what occurs while the learner constructs and revises it, even when an AI system no longer requires that representation as input.

5. Educational Implications and Research Agenda

The framework shifts attention from whether students should continue learning to draw toward a broader question: which intermediate representations help learners form, examine, and revise intentions before and during collaboration with AI? Sketching is one instance of a larger representational family. The same question can be asked of diagrams in science, notation in mathematics, models in engineering, maps in spatial reasoning, or structured outlines in writing.

This perspective also has implications for assessment. Final AI-generated artifacts can conceal much of the reasoning that preceded them. Intermediate representations can make parts of that process visible: what the learner noticed, what was specified, what remained uncertain, what changed after AI intervention, and which computational suggestions were rejected. Process evidence may therefore complement product evaluation in AI-supported curricula.

The conceptual framework generates several propositions suitable for future empirical investigation:

- P1. Externalizing an emerging creative intention before AI interaction may change the criteria by which learners evaluate subsequent AI outputs, compared with beginning from prompting alone.
- P2. Sketch-first and other representation-first workflows may produce different patterns of iteration from prompt-first workflows, including more movement among representation, generation, critique, and revision.
- P3. The degree and type of ambiguity preserved in an intermediate representation may influence whether AI contributions are experienced as productive exploration, unwanted default completion, or both.
- P4. Intermediate representations created before AI generation may support clearer accounts of the provenance of creative decisions and may strengthen learners' capacity to exercise reflective judgment over AI suggestions.

- P5. Productive friction introduced by representation-building activities may deepen observation and reflection even when it reduces the speed of artifact production.

Testing these propositions will require designs that distinguish observation from interpretation. Comparative studies could examine sketch-first and prompt-first conditions while collecting process traces, successive representations, reflective accounts, revision behavior, and measures of agency or ownership. Studies should also seek disconfirming cases: situations in which sketching adds little, constrains exploration, or is abandoned in favor of an AI-generated alternative. Such cases would help establish the boundary conditions of the framework rather than simply confirm it.

The framework is limited by its origin in reflective practice within the one elective art and design course. The instructor's simultaneous roles as teacher and framework developer necessarily shape the interpretation presented here, and students' elective participation may distinguish them from learners in other educational contexts. Experiences may also vary across creative disciplines and across generative systems with substantially different interaction capabilities. The framework has not been validated through controlled comparison, and the classroom experiences that motivated it should not be generalized to other learners, disciplines, technologies, or representational forms without further study.

6. Conclusions

Generative AI has made artifact production extraordinarily easy, but ease of production does not eliminate the cognitive work through which intentions form. It may instead make that work more important to identify and cultivate. This paper has proposed that sketching can function as an intermediate representation within human-AI creative cognition: a place where partially formed ideas become visible, remain revisable, and can be compared with computational interpretations.

The framework reframes sketching from a prerequisite for visual production to a cognitive technology for perception, reflection, communication, and judgment. Its significance lies not in opposing AI or preserving traditional technique for its own sake. Rather, sketching illustrates how an external representation can give human thought a form before AI expands it, creating opportunities for ambiguity, negotiation, revision, and continued agency.

The broader research question extends beyond drawing. As intelligent systems become collaborators across educational disciplines, researchers will need to understand how intermediate representations shape the relationship between human intention and computational generation. Researchers may need to revisit that understanding if computational systems are changing rapidly and are differentiating into specialized agentic services. The propositions offered here provide a starting point for that empirical work.

Abbreviations

The following abbreviation is used in this manuscript:

AI	Artificial Intelligence
UEAI	Understanding and Exploring Artificial Intelligence

References

(Baudoux et al., 2025) Baudoux, G., Guo, C., & Goucher-Lambert, K. (2025). Multimodal generative AI for conceptual design: enabling text-based and sketch-based human-AI conversations, *Proceedings of the Design Society*, 5, 2501-2510.

(Campbell, 2025) Campbell, B. (2025). Developing Collaborative Artificial Intelligence in Artistic Practices: Using AI in Creative Explorations, *IEEE Computer Graphics and Applications,* 46(1), 123-131.

(Darvishi et al., 2024) Darvishi, A., Khosravi, H., Sadiq, S., Gašević, D., & Siemens, G. (2024). Impact of AI assistance on student agency. *Computers & Education*, 210, Article 104967.

(Edwards et al., 2024) Edwards, K., Man, B., & Ahmed, F. (2024). Sketch2Prototype: rapid conceptual design exploration and prototyping with generative AI, *Proceedings of the Design Society*, 4, 1989-1998.

(Goel, 1995) Goel, V. (1995). *Sketches of Thought*. MIT Press.

(Goldschmidt, 1991) Goldschmidt, G. (1991). The dialectics of sketching. *Creativity research journal*, 4(2), 123-143.

(Guo et al., 2025) Guo, W., Liang, Z., Wang, C., Li, X., Hu, H., Chen, S., Yu, Q., & Zhao, Q. (2025). Student-AI collaborative creative problem-solving: the role of human agency. *Computers & Education*, 239, Article 105433.

(Hutchins, 1996) Hutchins, E. (1996). *Cognition in the Wild*. MIT Press.

(Kirsh, 2010) Kirsh, D. (2010). Thinking with external representations. *AI & Society,* 25(4), 441–454.

(Knoth et al., 2024) Knoth, N., Tolzin, A., Janson, A., & Leimeister, J. M. (2024). AI literacy and its implications for prompt engineering strategies. *Computers and Education: Artificial Intelligence*, 6, Article 100225.

(Oppenlaender et al., 2023) Oppenlaender, J., Linder, R., & Silvennoinen, J. (2023). Prompting AI art: An investigation into the creative skill of prompt engineering. *International Journal of Human–Computer Interaction,* 41(16), 10207–10229.

(Paananen et al., 2024) Paananen, V., Oppenlaender, J., & Visuri, A. (2024). Using text-to-image generation for architectural design ideation. *International Journal of Architectural Computing*, 22(3), 458–474.

(Pea, 1993) Pea, R. (1993). A Heuristic Framework for Raising Empirical Questions on Distributed Intelligence – Artifacts as Cognitive Offloading, In G. Salomon, (Ed.), Distributed Cognition: Psychological and educational considerations (pp. 47–87). Cambridge University Press.

(Rafner et al., 2025) Rafner, J., Zana, B., Hansen, B, et al. (2025). Agency in human-AI collaboration for image generation and creative writing: preliminary insights from think-aloud protocols, *Creativity Research Journal*, Advance online publication.

(Rezwana & Maher, 2023) Rezwana, J. & Maher, M (2023). Designing Creative AI Partners with COFI: A Framework for Modeling Interaction in Human-AI Co-Creative Systems, *ACM Transactions on Computer-Human Interaction*. 30(5), 1-28.

(Schön, 1983) Schön, D. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.

(Star & Griesemer, 1989) Star, S. and Griesemer, J. (1989). Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals. In Berkeley's Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science.* 19(3), 387–420.

(Suwa & Tversky, 2009) Suwa, M. & Tversky, B. (2009). Thinking with sketches. In A. B. Markman & K. L. Wood (Eds.), *Tools for innovation: The science behind the practical methods that drive new ideas* (pp. 75–84). Oxford University Press.

(Varela et al., 1992) Varela, F., Thompson, E., & Rosch, E. (1992). *The Embodied Mind: Cognitive Science and Human Experience*. The MIT Press.

(Yang et al., 2026) Yang, Y., Chen, L., Sun, D., Zhu, W., & Pang, S. (2026). Fostering preservice teachers' design thinking by integrating generative AI and analytics-supported reflective assessment in a collaborative inquiry environment, *Computers & Education*, 249, Article 105602.