Tracing Non-Dyadic Creative AI Interaction

Carey Linda Crooks
Herding CATs: Making Sense of Creative Activity Traces CHI 2026 Workshop
studio@careylindacrooks.com

ABSTRACT

AI-based creativity support tool (CST) design often reflects a foundational assumption: that the unit of human-AI creative interaction is 1:1, or, one person, one tool. In fact, LLM integrations introduce an opaque, collective dimension. Models are trained on vast, heterogeneous data, embed many creative norms, and are further shaped by usage at scale. Despite this complexity, most interface design presents as a neutral, single interface. Current CAT methods employed with AI-based CST tend to solely map the journey of the human user, collapsing these details into a single trace. We propose that making this complexity more visible through CAT analysis is a productive provocation with the potential to reimagine CST interaction paradigms in the service of the creative process.

CCS Concepts: • Human-centered computing → Interaction Design, Human Computer Interaction (HCI), Visualization;

• Computing methodologies → Human-AI interaction

Keywords: creative activity trace (CAT), creativity support tool (CST), AI-based CST, non-dyadic

1 Reaching past the 1:1 assumption

Creative activity trace (CAT) analysis methods such as linkography typically map one person's creative process, perhaps while utilizing a creative support tool (CST). Techniques such as "fuzzy linkography" utilize the capabilities of machine learning to create linkages at scale. This 1:1 discrete framing is productive: it enables a controlled and specific environment for observation, especially as coding the traces can be manually intensive. [1] But it also encodes an assumption that deserves scrutiny as AI-based CSTs become more integrated in creative practice: that the human-AI creative interaction is fundamentally dyadic. Approaches toward "non-dyadic" interaction research have been well documented and explored in Human-Robot Interaction (HRI) space. However, while the HRI field describes configurations of humans and *embodied robots*, this paper uses the term non-dyadic to describe the multitudinous composition of a seemingly singular, non-embodied CST interface.

An AI-based CST is not a discrete instrument. Utilized in a creative context, it is a compression of immense training data containing aesthetic traditions and design decisions made across many contexts. The model's suggestions reflect the creative norms of the communities represented in its training data. Yet, often in the form of a chat interface, the model behaves as a single, consistent interlocutor. When a practitioner incorporates an AI-based CST into their practice, and utilizes an output, they are not just making a creative decision but they are in dialogue with an aggregate of prior decisions. This dialogue extends not only to past creative choices embedded in the data, but also to the broader concurrent user community. Current CAT methods do not trace this broader interaction, focusing instead on the individual human user's experience, or point of view, of the interaction.

We propose that visualizing these interwoven processes would be an advancement for the field of creativity and AI-based CST. How can CAT analysis methods, and the nature of AI-based CSTs, be extended to capture the collective, multi-user dimensions of creative process? Answering this call requires moving from a 1:1 interaction model toward two alternative framings: *1:many*, in which a single creative practitioner is understood to be in dialogue with an expansive embedding; and *many:many*, in which communities engage with AI systems that are themselves aggregates of many communities' creative practices. Each framing demands different trace methods and may inspire different CST designs. If CAT analysis methods allow us to peer into the creative process, how might they be deployed to visualize a larger creative system, involving LLM chain of reasoning, referential data corpus, and the live human creative process?

2 Research Directions

2.1 Tracing the 1:Many: Situating the Practitioner in Dialogue

The first direction addresses the need to peer beyond the neutral façade of AI-based CSTs. While HRI is mature in frameworks for analyzing multi-user and robot interactions, it is trickier to define interaction with non-embodied AI-based CST tools. When practitioners interact with an AI-based CST and accept, modify, or reject AI suggestions, their decisions are enmeshed with a broader system with limited transparency. In creative use cases, CST output is presented without indicating

what data it is drawing on, despite this output being shaped by the work of other practitioners, past and present. Advancements in LLM transparency like Chain of thought (CoT) [3] can be useful to the work of a broader human-AI CAT analysis. However, they stop short of a visualized interaction paradigm. Recent work like Graph of Thought (GoT) [4] gets closer to this aim. Beyond step-by-step reasoning, attributing reasoning to specific sources, such as via provenance, helps make the model's collective constitution visible rather than masking it behind a singular interface. CAT analysis should draw on these advancements in model self-reporting to visualize the fuller picture, to better illustrate the weaving of human design moves alongside machine influence.

2.2 Tracing the Many: Many: Reimagining the CST as an Interface between Amassed Processes

AI-based CSTs are constantly evolving tools and not just via initial bulk training, but through further, innovative loops of refinement such as reinforcement learning from human feedback (RLHF) [5]. It is an interesting design problem: how to make these loops of user input part of the 1:many tool experience? For instance, what would it look like for a user to better understand the relationship between their creative moves and those of others? Or to try on a new creative process explicitly? Bringing creative process information to the interface level allows creatives to negotiate with the processes they choose to use and understand the active role they play in the processes available, or reinforced, for others.

Additionally, newer autonomous architectures, i.e. agents, align with this framing of many: many, and further complicate visibility into decision-making and need to be accounted for via CAT analysis.

2.3 Possible Applications in Designing for the Expanded Model

The new tracing conventions formed by these directions have application not only in meta-analysis of creative processes but also for informing the design of the CST interface itself. These changes may take the form of reimagining the CST to no longer be a design *entity* but an interface to a plethora of design processes. We believe that surfacing this information can have profound benefits for the individual's creative process and decision-making.

In the scope of this proposed research, interaction with AI-based CSTs differs physically from those of embodied robots. This assumption may need to evolve if product form factors change, become embodied, and the line between robot and AI-based CST blurs.

3 CONCLUSION

The 1:1 model of human-AI creative interaction limits the research potential of CSTs in the field of creativity. Visualizing human creative moves alongside AI reasoning in the trace record requires new methods, new coding conventions, and new theoretical frameworks, all of which could have profound positive impact in CST design itself. Researching tracing conventions for 1:many and many:many interaction models is the research agenda we propose.

REFERENCES

[1] Amy Smith, Barrett R Anderson, Jasmine Tan Otto, Isaac Karth, Yuqian Sun, John Joon Young Chung, Melissa Roemmele, and Max Kreminski. 2025. Fuzzy Linkography: Automatic Graphical Summarization of Creative Activity Traces. In Proceedings of the 2025 Conference on Creativity and Cognition (C&C '25). Association for Computing Machinery, New York, NY, USA, 637–650. https://doi.org/10.1145/3698061.3726915

[2] Eike Schneiders, EunJeong Cheon, Jesper Kjeldskov, Matthias Rehm, and Mikael B. Skov. 2022. Non-Dyadic Interaction: A Literature Review of 15 Years of Human-Robot Interaction Conference Publications. J. Hum.-Robot Interact. 11, 2, Article 13 (June 2022), 32 pages. https://doi.org/10.1145/3488242

[3] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In Proceedings of the 36th International Conference on Neural Information Processing Systems (NIPS '22). Curran Associates Inc., Red Hook, NY, USA, Article 1800, 24824–24837.

[4] Yao Yao, Zuchao Li, Hai Zhao. 2024. Beyond Chain-of-Thought: Effective Graph-of-Thought Reasoning in Language Models. arXiv:2305.16582. Retrieved from: https://arxiv.org/abs/2305.16582

[5] Chip Huyen. 2023. RLHF: Reinforcement Learning from Human Feedback. huyenchip.com. Available at: https://huyenchip.com/2023/05/02/rlhf.html