

An Account of Contribution Dissolution in a Research Team^{*}

Jared Ren

renjared@uw.edu

Human Centered Design & Engineering, University of Washington
Seattle, USA

Abstract

In spring 2026, our four-person graduate research team spent a quarter interviewing professional visual artists about how they defend authorship when generative AI enters their work. The artists had clear rules about where AI was allowed. We had none, and by the end of the quarter a word none of us could account for had traveled from a coding memo into our final presentation, where our advisor asked why we had chosen it. Nobody could say. This is a first-person account of that collaboration and its quiet breakdowns: coding decisions their own authors could not explain, a suspicion I never voiced, a flagging rule whose silence carried no information, and a disclosure statement that became true only after we had retreated to working manually. I argue that AI does not need to write anyone’s work to dissolve contribution; it only needs to be available in the background of each member’s individual workflow. The conditions that make collaboration accountable live in people, not in records, and no record we kept, or could have kept, would have restored them.

CCS Concepts

• **Human-centered computing** → **Empirical studies in collaborative and social computing.**

Keywords

contribution dissolution, accountability, generative AI, AI disclosure, collaborative work, qualitative research, autoethnography

1 Introduction

In spring 2026, four graduate researchers, of whom I was one, conducted a qualitative interview study of how professional visual artists maintain creative identity and authorship when generative AI enters their workflow. The artists we interviewed had clear, principled rules about where AI was allowed in their work: text in, images out; keep execution manual; use AI as reference, never as final work. They could say exactly where AI was permitted in their process and defend the reasons. We admired the clarity of those rules. We had nothing like them ourselves.

What follows is a first-person account, reconstructed from memory, team documents, and records of my own AI use. My three teammates appear anonymized, and I want to be explicit about a limitation that is also my argument, in that much of what I report about their AI use is inference. I could not reliably tell, during the project or now, where their judgment ended and a model’s output began. Neither, I will suggest, could they always tell about themselves.

^{*}Position paper accepted to the workshop “Nobody Did This”: Contribution, Originality, and Accountability in Agent-Mediated Collaboration [5] at the 2026 ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW ’26), Salt Lake City, UT, USA, October 2026.

Table 1: The six conditions of accountable collaboration that the scenes show breaking.

	Condition
C1	Every member understands the shared goal.
C2	Every member understands their own work and, in outline, everyone else’s.
C3	Every member can tell their intentional from unintentional choices, and can give reasons on demand.
C4	Every member knows that the others meet C1–C3 too.
C5	Every member knows who to ask, and trusts the answer will be good.
C6	Questions asked to the right member get answered well.

The claim this case supports: AI does not need to write anyone’s work to dissolve contribution, rather, its mere presence in the background of each member’s individual workflow is sufficient. What breaks is a set of social conditions that never lived in any record, so no record can restore them.

2 Four scenes

The scenes below are ordered by depth. The first breaks the work, the second breaks the text, the third breaks trust, and the fourth breaks the record. Figure 1 maps our workflow and where AI entered it; Table 1 names the six conditions of accountable collaboration (C1–C6) that the scenes will show breaking.

2.1 The codes

The collaboration began well. We formed research questions together, in person and over video. We drafted interview questions together. Every interview had at least two members present, a moderator and a notetaker. Through this phase, each of us could see the others’ contributions being formed, and could have answered for our own.

Analysis was where the work became individual, and where AI entered. For first-round coding we each took one transcript. Everyone used AI, but differently: one member ran a transcript through a model for a complete first pass; another asked a model how one ought to code before coding themselves. Our first-round methods diverged, and that was expected. Human coders diverge; consolidation meetings exist to work over the differences. We consolidated, built a codebook in team sessions, and the codebook was ours.

Second-round coding broke differently. Each member coded a new transcript against the shared codebook, again with AI, now in ways that varied and went undisclosed. Some uploaded the codebook and the full transcript and worked from the model’s output. Some consulted the model only on confusing passages. When we for a team discussion and someone asked the ordinary question of




Figure 1: The team’s workflow, left to right. Collective phases are light; individual phases are dark. Labels beneath a phase are AI uses that entered there, colored by outcome: filled red, conditions broke; outlined red, strained but repairable; blue, conditions held. Parenthetical tags list the conditions each use broke or strained (Table 1).

collaborative analysis, “why did you code it this way?”, there was no clear answer. No one was being defensive, and no one disagreed. The reasons were simply absent. The person who had done the work was no longer the authority on it. We spent extra sessions retracing steps, redoing coding, and reassigning work that could not be salvaged.

Two mechanisms fed this. The first was deferred authorship. More than once, someone would say some version of “I used AI here, but I’ll work on it myself later; this is just initial drafting. I was busy this week.” Later rarely came, because waiting on one member’s revision would have stalled everyone, so the team moved forward and the temporary AI output became permanent. The second was trust itself. I did not have the time or the standing to check everything my teammates produced; deadlines existed, and a baseline of trust is what makes teamwork possible at all. Through that trust, AI-mediated work entered our shared documents unmarked.

What broke was answerability, the capacity to give reasons on demand (C3). Divergence between human coders arrives attached to reasons, and reasons are what a consolidation meeting works with. AI-mediated divergence arrived without reasons, and there was nothing to consolidate.

2.2 The word

In one of our comparison memos, a section heading read “Sarcasm as the Frontier AI Cannot Cross,” using frontier to mean a limit, the thing AI could not do. By our final presentation, the word had changed. A passage describing the limits artists set on AI now read, in language close to this: “that line tracks AI’s capability frontier, and the frontier is moving. These aren’t permanent settlements; they’re live negotiations that will have to be redrawn.” The limit had become an expansion: frontier, settlements, territory to be redrawn.

After the presentation, our advisor flagged the language as “colonial vocabulary, in a project about artists defending their work from extraction.” They asked why we had used those words. I was

confused; I did not recognize the choice as my own so I handed the question to my teammate who owned the section. They could not answer. My own rule, throughout the project, was to check anything AI touched until I could explain it, so I knew what that silence meant, and in it my suspicion finally took a definite shape.

Here is what should be unsettling: we had the paper trail. I found the word’s first appearance in our memos. However, none of it answered the question our advisor actually asked: who chose it, and why. The record showed where, but it could not show where it came from. A word choice, and the worldview that came with it, had entered our work without ever being anyone’s decision (C2–C4).

This is the workshop’s title made literal. Somebody’s process produced that sentence. Nobody did it.

At the next team meeting I proposed a rule: if you create something with AI, make a note so we can flag it before it sneaks into our shared work. Everyone agreed. Less formally, we retreated to coding by hand, drafting in our own words, and redoing what could not be accounted for. It was the only way we could trust the work again.

2.3 The question

When we exchanged analytic memos, I read one from a teammate and could not shake the sense that I was not reading them. The writing was full of confident interpretation, one staged reveal after another, one “it’s x, not y” construction after another. None of these are proof of anything; the same features appear in plenty of human writing, and studies show people are bad at judging what is AI-generated [6]. What unsettled me was narrower: the gap between this written voice and the person I talked with every week. When we later asked them to walk the team through the memo, they said there wasn’t much to say. The memo was rich but its author had little to add. The document held more interpretation than its author could speak to (C3–C6).

I said nothing. Instead, I reasoned that they were busy, that the workload had been heavy that week, that they were doing their best against their own deadlines. I also, frankly, did not want to embarrass a teammate I liked and worked well with, in a small team whose energy we all depended on. To say “this reads as AI-generated, can you redo it” is an accusation no matter how it is meant, and one I might have been wrong about. So instead of asking, I absorbed the cost: quietly reworking, cross-checking against transcripts, and eventually, in our final document, making more than a hundred sentence-level edits to remove vocabulary that was alien to the language we had agreed on as a group.

Call this the unaskable question. The ordinary repair mechanism of collaboration, simply asking a colleague to account for their work, now risked insulting them, and staying silent made sense for everyone in the moment. There was even a reason for teammates to stay quiet that had nothing to do with bad faith: people consistently value work less once it is labeled AI-assisted [7], so disclosing its use carried a real penalty in how one’s work would be received. I did not ask and they did not tell, which led to my trust in the team wearing down, invisibly, without a single conversation in which it could have been repaired.

2.4 The disclosure

Our course encouraged an AI-use disclosure from the start of the quarter. Through the months when AI was actually woven through our coding and drafting, we never wrote one. The flagging rule I proposed mid-project did not work much better: everyone agreed, and flags rarely appeared, which left me unable to interpret the silence. Did no flag mean no AI, or just nobody following the rule? A disclosure rule where staying silent looks the same as not using AI tells you nothing at all.

The one disclosure we did write appears in our final report:

AI-use disclosure. We used AI tools during drafting and revision: generating draft language for some sections, exploring framing alternatives, and checking quotes and terminology for consistency. The research itself is ours: we conducted and coded the interviews, built the codebook in team sessions, and made all analytic and structural decisions in team discussion. Every quote in the paper was verified against the original transcripts by us, and all AI-suggested text was reviewed, revised, or rejected by the team before it entered the document. We take full responsibility for the content.

By the time we wrote this, it was true. That is what should trouble us. It was true because of the sequence the previous scenes describe: it came after the retreat to manual coding, after the redone work, after my edits removing the alien vocabulary. In the report’s production, AI had been limited to generating templates and suggestions on top of writing that was already ours. Every claim in the statement had, by then, a practice behind it. That final, describable role is the blue label at the far right of Figure 1. For the months the statement does not cover, when AI was woven through the analysis itself, no honest version of it could have been written, and none was.

Call this disclosure by retreat. The disclosure did not make our process accountable; abandoning the process made the disclosure

writable. This, in miniature, is the trap in our current safeguards. Watermarks and detectors assume fully AI-written output, and reliable detection may not even be possible [8]. Disclosure statements assume that the facts needed for attribution exist and merely need reporting. In our team those facts were never produced. None of us was careless. AI entered at the level of individual sentences and individual codes, smaller than the level at which any of us notice our own choices.

3 What broke

A skeptic could brush these scenes off: novice coders often cannot explain their coding decisions, and course teams have always had coordination problems. The case itself answers this. The same four people, in the same quarter, produced fully answerable work in every phase done collectively, forming questions, interviewing in pairs, consolidating the codebook, and unanswerable work in the phases where AI sat in the background with each member alone. The pattern is visible at a glance in Figure 1: every AI use that broke or strained a condition sits under an individual phase. Both breakdowns surfaced in collective settings, the coding discussion and the presentation. The failure also had a specific shape: not bad reasons, which novices produce and meetings repair, but the absence of reasons, authors with nothing to say about their own richest passages.

Collected in Table 1, the conditions the scenes tagged one by one form a system, and the order in which it failed matters. C3 broke first, and it broke inside each person: judgment was handed to the model in ways that left no reasons behind, so the person who did the work could not answer even for themselves. This is what I think is new compared to the breakdowns CSCW has long studied. In collaboration, mutual knowledge fails between people who each know their own work [3]; here, common ground [2] failed at its source, because there was no self-knowledge for the team to share. The damage then spread outward. Once one “why” went unanswered, each of us privately lost confidence that the others knew their own work, and knew the others had lost that confidence in us (C4). Doing the work no longer made you the person who could answer for it, so there was no longer a right person to ask (C5, C6), and the team’s map of who knows what failed with it [10]. Through all of this, everything looked normal from the outside: deadlines met, workload evenly split, meetings friendly. That is why no awareness mechanism [4] flagged anything, and why the articulation work of keeping visible who is doing what [9] continued smoothly while the thing it tracked ceased to exist.

There is a second pattern the figure does not draw. Our uses of AI differed in how much was handed to the model and in how visible the use was to the team, and the two mattered differently: the more that was handed over, the worse the damage. The more visible the use of AI, the easier the repair. Two cases show the two effects apart. Consulting the model only when confused was nearly invisible but delegated little, and the damage stayed contained. The disclosed AI draft was heavily delegated but announced at its creation, and the team could work around it. None of our uses was “fully AI,” which is why detection-style safeguards could not have caught any of it.

4 What we lost

The artists we interviewed could defend their authorship because their medium made it visible where AI entered. An image is generated or it is not; execution is manual or it is not; a rule like “text in, images out” can be stated, followed, and seen to be followed. In collaborative knowledge work, AI enters at the level of sentences, codes, and word choices, where even the person using it cannot see the line.

When our collaboration finally repaired itself, it did so by stumbling into a rule of its own about where AI was allowed. This appeared in AI on top of human writing, templates and suggestions on text we had already written, and nothing earlier in the analysis. Our disclosure became true once our practice became describable. But we should be honest about what that repair cost and where it worked. We were a course team with redoable work and low stakes since we could afford to redo a quarter’s coding by hand. Professional teams facing production deadlines cannot simply retreat. Our fix says more about how few options exist than about what anyone should do.

We did not lose the record. The memos survived and the word’s entry point was findable. No tool could have recovered the missing fact, who exercised judgment where, because that fact was never created in the first place. What we lost was witnessing the work, where we are working in ways that make our choices answerable to each other, and knowing that we all do [1]. For at least one word we presented, “nobody did this” is the literal truth.

So I end with the question our case could not answer. The artists had rules about where AI was allowed because, for them, that placement was visible. What would such a rule even look like in collaborative knowledge work, where AI enters at a level invisible even to the person using it, where asking about it is an accusation, and where disclosing it carries a penalty? Until we can answer that, every disclosure statement in every paper describes a policy and witnesses nothing.

References

- [1] Mark S. Ackerman. 2000. The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility. *Human-Computer Interaction* 15, 2–3 (2000), 179–203. doi:10.1207/S15327051HCI1523_5
- [2] Herbert H. Clark and Susan E. Brennan. 1991. Grounding in Communication. In *Perspectives on Socially Shared Cognition*, Lauren B. Resnick, John M. Levine, and Stephanie D. Teasley (Eds.). American Psychological Association, Washington, DC, 127–149.
- [3] Catherine Durnell Cramton. 2001. The Mutual Knowledge Problem and Its Consequences for Dispersed Collaboration. *Organization Science* 12, 3 (2001), 346–371. doi:10.1287/orsc.12.3.346.10098
- [4] Paul Dourish and Victoria Bellotti. 1992. Awareness and Coordination in Shared Workspaces. In *Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work (CSCW '92)*. ACM, 107–114. doi:10.1145/143457.143468
- [5] Kashif Imteyaz, Mohammad Rashidujjaman Rifat, Divya Ramesh, Steven R. Rick, Simo Hosio, Hauke Sandhaus, Advait Sarkar, Christoph Riedl, and Saiph Savage. 2026. “Nobody Did This”: Contribution, Originality, and Accountability in Agent-Mediated Collaboration. arXiv:2607.26387 [cs.CY] doi:10.48550/arXiv.2607.26387 Workshop at the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW '26).
- [6] Maurice Jakesch, Jeffrey T. Hancock, and Mor Naaman. 2023. Human Heuristics for AI-Generated Language Are Flawed. *Proceedings of the National Academy of Sciences* 120, 11 (2023), e2208839120. doi:10.1073/pnas.2208839120
- [7] Manav Raj, Justin M. Berg, and Robert Seamans. 2026. The Artificial Intelligence Disclosure Penalty: Humans Persistently Devalue AI-Generated Creative Writing. *Journal of Experimental Psychology: General* (2026). doi:10.1037/xge0001889 Advance online publication.
- [8] Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. 2023. Can AI-Generated Text Be Reliably Detected? arXiv:2303.11156 doi:10.48550/arXiv.2303.11156
- [9] Kjeld Schmidt and Liam Bannon. 1992. Taking CSCW Seriously: Supporting Articulation Work. *Computer Supported Cooperative Work (CSCW)* 1, 1–2 (1992), 7–40. doi:10.1007/BF00752449
- [10] Daniel M. Wegner. 1987. Transactive Memory: A Contemporary Analysis of the Group Mind. In *Theories of Group Behavior*, Brian Mullen and George R. Goethals (Eds.). Springer, New York, NY, 185–208.