

“I Think We Did This”: GenAI Contributes to Erosion of Ownership and Design Intention in Group Design Work*

Hauke Sandhaus
hgs52@cornell.edu
Cornell Tech
New York, NY, USA

Ruixiang Han
rh652@cornell.edu
Cornell Tech
New York, NY, USA

Wendy Ju
wendyju@cornell.edu
Cornell Tech
New York, NY, USA

ABSTRACT

Increasingly widespread generative AI adoption is changing the nature of group design work. To understand how ownership and design intention shift, we interviewed eighteen students from a graduate interactive device design course where projects are done in teams, each walking through their own repository and chat logs, and we report an emerging finding in three parts. First, design intentions can have histories their authors cannot recover. Final projects narrated as personal ideas traced back to AI brainstorms the students could not place or did not recall. Second, intention forms downstream of the tool. Students discover their preferences by selecting among generated options, ask the model to argue for its own suggestions, and accept its case for scaling an ambition down. Third, the group record dissolves. The model merges teammates’ contributions, its dialogues live in single members’ private accounts, and students turn to *design intent cosmetics* such as re-drawing generated storyboards by hand, work done after the fact that restores a feeling of ownership while obscuring the actual process. We read these accounts through Lessig’s pathetic dot, the designer shaped by an architecture they no longer see. In the worst cases, the designer’s authority over the work erodes, both their intention and their ownership.

KEYWORDS

generative AI, group design work, ownership, attribution, accountability, design education, agent-mediated collaboration

1 POSITION: THE PATHETIC DESIGNER

Our position is that generative AI is eroding design students’ authority over their own work, both their intention and their ownership, and that group design work is where the erosion is hardest to see. We illustrate the position with preliminary evidence from student interviews. In these episodes a final concept traced back to an AI brainstorm its author could not recall, a documented design process ran backwards, and students applied *design intent cosmetics* after the fact to make the work feel owned.

Lessig’s account of regulation turns on an unflattering figure he calls the *pathetic dot*, the individual acted on by several forces at once and shaped most thoroughly by the one least visible to it, the architecture of the systems it acts through [14]. The dot is pathetic because it is regulated from every side and may take the result for its own free choice.

The designer is our discipline’s counter-image to the dot, the person who makes architecture rather than being made by it and whose defining competence is holding an intention and imposing it on material. Design education exists to produce such people. We observe design students, not working professionals, but design education is where the counter-image is supposed to be formed, so what happens there matters for the discipline. We wrote this position paper for the CSCW 2026 workshop “Nobody Did This” [11], which asks who did what, and who is accountable, when agents mediate collaboration. Our impression from watching these novice designers is that in AI-saturated design work the designer can come to occupy the dot’s position, shaped by the architecture of the tools they work through without being aware of it, and that group design work is where this is hardest to see and matters most.

Our position grew out of teaching and grading student design teams in our graduate interactive device design course. We began to notice odd sloppiness in student design documentation, signs we later learned to read as tell-tale marks of AI-generated text. Over time the documentation came to look better, perhaps because students gained proficiency with the tools. Two worries remained. The first is that we no longer know whether, or what, students still learn, when they have delegated different aspects of the design work to AI tools. The second is that the traditional indicators of good student design work have stopped being reliable. Completeness of documentation is the simplest of these. Templates come back filled, and paragraphs of text describe thought processes that the students have not thought. The interviews below show how this happens from the inside. What the grader can no longer see is careful design work. Some of the sloppiest-looking submissions turned out in interview to rest on laborious, deliberate work, invisible among generated material that looks like everyone else’s. And some of the most careful students kept generated text out of their submissions entirely, so that their own work would remain visible. A course that cannot tell careful work from generated output grades the wrong thing.

The tool’s architecture now takes part in forming the intention itself. It suggests the concept, surfaces the preference, and fills the gaps between teammates, and this participation leaves no trace the designer can later find. Students still present themselves as authors, in a collective first person that includes the model without saying so. The narration is a reconstruction, and once the design history is traced back, the record sometimes disagrees. What erodes is the designer’s authority, their design intention and their ownership of the work, and the practices below reach both.

The fear that the humans, not the machines, become the automata is older than this technology. Marx wrote of the worker as

*Position paper accepted to the workshop “Nobody Did This”: Contribution, Originality, and Accountability in Agent-Mediated Collaboration [11] at the 2026 ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW ’26), Salt Lake City, UT, USA, October 2026.

"an appendage of the machine" [18], Winner of the *reverse adaptation* of human ends to the available means [27], Anders of a *Promethean shame* before the perfection of one's own devices [3], and Bainbridge of the operator left whatever automation could not absorb [4]. The posture is recognizable in our students' accounts as steady production without recoverable intention. On the agency spectrum recent CSCW work uses for human–AI co-creation [31], these accounts fall at the passive end at the two moments that matter most, concept formation and integration.

HCI has studied most of the pieces of this picture separately, and at the level of the individual. Designers have learned to treat machine learning, and now generative models, as a design material [5, 7, 29], and practitioners report gains alongside worries about originality and deskilling [15, 16, 20]. Steering the model is hard for non-experts [30]. Knowledge workers who lean on generative AI report less critical thinking and less confidence in their own judgement [8, 13]. On ownership, Draxler et al. [6] found that people who did not feel ownership of generated text still declared themselves its authors, and Xia et al. [28] found workplaces in which hiding AI use well counts as skill. The responses run through declaration, from the ACM's authorship policy to disclosure schemes [1, 2], and design education is rebuilding its policies and assessments around the same tools [10, 12, 17].

The group side is thinner. Star and Strauss [25] showed that what a system keeps account of decides which work stays visible, and the work on generative AI in groups has so far asked how people work with the tool together. Teams that adopt generative AI develop shared prompting strategies [9], AI has been placed inside group ideation [24], and data scientists described automated AI as a collaborator rather than a tool [26]. The earlier study of this course interviewed twelve project groups and found the gains concentrated in execution [22]. The workshop this paper is written for names the resulting problem contribution dissolution [11]. All of this work assumes that contributions come from members and are witnessed by members, and the declaration schemes assume the same. An agent that merges, drafts, and remembers inside the workflow removes the witness, and the accounts below show what the record looks like once it is gone.

Some students in the same corpus govern these tools well. They specify before generating, bound the interface, and refuse generic defaults. This paper is not about them. Our evidence is retrospective self-report anchored by artifacts, from one course, gathered by the people who taught it. We report the cases where generative AI has evidently contributed to erosion of ownership and design intention in student group design work.

2 EVIDENCE OF EROSION

2.1 Study and method

We report preliminary evidence collected from *Interactive Device Design*, a master's-level, project-based prototyping course at Cornell Tech. Students build functional physical–digital devices across a sequence of two-week lab sprints and a self-directed final project, working alone or in teams of two to four. In the Fall 2025 offering studied here, teams built on Raspberry Pi 5 with sensors, actuators,

and AI APIs, and generative AI use was permitted throughout and part of everyday practice.¹

We interviewed eighteen students (P1–P18), individually, after final grades were released, in semi-structured sessions of roughly 30–60 minutes (Cornell IRB0148132, exempt). As an incentive, participants could enter a \$50 raffle. Each interview walked the student through their own repository, lab by lab, with their AI chat logs open where they survived, prompting recall of specific decisions rather than general impressions. The interviews were conducted by members of the teaching team, so students already knew their interviewer from the course, and with grades already released, candor carried no academic risk. Recordings were transcribed and anonymized. Transcripts were coded inductively by three independent automated coders with researcher adjudication at fixed checkpoints, yielding a 91-code system organized by Pugh's account of autonomy [21]. Quotes are verbatim, trimmed for length only. A minority of interviews are documented by detailed contemporaneous notes rather than full transcripts, where a participant declined audio recording or a recording was lost. Prior publications from this course report an earlier cohort under a different question [22] and repository telemetry across three cohorts [19], and the professional-autonomy framing is developed in a companion workshop [23]. This paper draws only on the 2025-cohort interviews.

2.2 Finding 1: Intentions with histories their authors cannot recover

A few instances of specific, direct erosion of design authority follow. The episode that anchors this paper (Figure 1, case 1) unfolded across one interview. Asked about the origin of their final project, a soothing "smart plushie," P7 gave a personal account:

"I just like plushies, and I always wanted to make something like a smart plushie. My teammate also has a lot of plushies in her house, and we just thought it's fun to do some projects around there." — P7

Asked whether the idea predated the final project, P7 said no: *"I think this is only for the final project"* (P7). Later in the same interview, P7 scrolled through an earlier sensors lab in the same repository, whose sketches turned out to have been brainstormed by ChatGPT from the team's parts list. The concept was among them:

"This positive thing, sensor thing, I think it was influenced by the ChatGPT... Okay, I'm gonna be very honest. I think while we're drawing these, we were considering which one we want to go as an actual idea, as a final project." — P7

The practice behind the episode is routine in this course, not one student's habit. Teams feed the model their parts list and select among its returns: *"I remember the idea is come up with me and AI together... I just gave the AI all of the sensors we have"* (P12). And we do not read the episode as dishonesty. P7 volunteered the correction as soon as the artifact made it visible, and elsewhere divided credit with care. The problem is harder than dishonesty. The

¹Course materials and lab structure: <https://github.com/IRL-CT/Interactive-Lab-Hub/tree/Fall2025>




Figure 1: Retracing an idea after AI-mediated group design work, the record and the origin story disagreed. At interview, students told the origin story of their final concept, then revisited the lab history with the interviewer (red). The trace tangles, and where it reaches an origin, the story gets revised. Cases from the corpus: (1) P7’s “personal” final concept was in an earlier lab’s AI brainstorm, found only by scrolling back; (2) P4’s assistant, with memory on, carried an AI-brainstormed idea forward until it felt self-generated; (3) P11 could not place which lab an idea came from; (4) P12 could not say whether an image in the team’s documentation was AI-generated or a teammate’s.

influence was invisible in the telling and recoverable only through the repository and chat history, and had the interview not happened, no one, including P7, would have known. Other interviews show the same pattern, less sharply. P7’s own review practice concedes it: “I usually double-check with, like, whatever the assignment needs, like, I’m like, did I write it? Did I write them? And then... I will ask my teammate” (P7). Authorship of one’s own submission is something to look up. Another team’s game concept arrived in an AI brainstorm. The interview notes record that the team had no good idea first, asked AI, and felt inspired, and the same participant later credited the ideas as 70 percent the team’s own (P4), a split at odds with the notes. The channel for the influence can itself be invisible: “I just didn’t turn it off. So my memory feature is always on in my ChatGPT” (P4), carrying material between projects and shaping later brainstorms with context the student no longer thinks about. A third origin story ran the same way. P11 introduced the final project as personal, “I’m just a big subway enthusiast” (P11), and, asked whether it began earlier, could not place it: “there was some inkling of something I’d worked on already for a previous lab assignment... I’m honestly forgetting which lab that might have been” (P11). The same participant frames feeding a whole lab brief into the model as the exception to a sense of responsibility for the assignments: “I have done that in a lot of scenarios, copy the whole lab and throw it in chat” (P11). A chain from lab brief to model to lab submission to final project exists, but no one can follow it back.

2.3 Finding 2: Intention forms downstream of the tool

A second cluster of evidence shows intention forming downstream of the tool. Rather than an instrument picked up to execute a plan, the model shapes what gets made along the whole pipeline. Students discovered what they wanted by selecting from what it offered. The counterfactual was sometimes explicit: “if we didn’t try to

figure out more fun ideas, I probably would have done something really generic” (P16). The model gets credit for the work and for its distinctiveness. The adoption decision itself can run through the tool. One participant has the model argue for its own ideas and choose between them: “I kind of give it back both of those, and I’m like, oh, which one is better to do?” (P17). The idea, the test it passes, and the reasons for choosing it come from the same source, and the idea is then felt as one’s own. The negotiation also runs downward. P18 wanted a sign-language system that handled words, and the model advised against it, “the LLM basically said, this is going to take a very long time... I recommend you do letters instead” (P18). P18 built the letter version. Asked whether the model’s version won because it was better, “it’s faster and more convenient for me to figure out how I need to do something, especially if I don’t know how to do it” (P18). Where the loop ran unattended, students describe mechanized motion:

“It prevents you from taking the reins... and then by the end, you realize, like, oh shit, I covered so much ground, I’ve created so many artifacts and deliverables, and so much insights from it, without having to finally see where I’m heading towards.” — P5

“The more we tried to vibe code to fix it, the more we didn’t really understand what was going on.” — P6

The delegation reaches the artifacts the curriculum uses to teach reflection. Asked about the interaction diagrams in their repository, the course’s device for making a design intention inspectable, one participant looked at their own submission and said:

“This is a Verplank diagram. I don’t think I ever made a Verplank diagram of my own. I think most of them have been AI-generated.” — P17

AI also interferes with the methods whose purpose is design thinking itself. The course teaches Wizard of Oz so an interaction can be examined before it is built, and wireframing so a screen exists

Table 1: Practices that make contribution unreconstructible. Each is reasonable on its own, and together they remove, replace, or never create the evidence any accountability scheme would need.

Practice	Seen in	What the record loses
<i>Removing the record</i>		
Removing AI’s style markers from text	P11	the surface signal that text was generated
Deleting chat logs	P5	the only process record
<i>Replacing the record</i>		
Redrawing generated artifacts by hand	P15	the line between made and adopted
Splitting credit retroactively, from memory	P11, P16	substitutes arithmetic on memory for a record that does not exist
Process artifacts generated after the fact	P11	the record now shows a process that did not happen
AI supplies the idea and then the case for adopting it	P17	any ground for the choice independent of its source
<i>A record no one holds</i>		
Skipping the method steps that would show intention forming; building the artifact directly	P5, P11	the intermediate artifacts (sketches, wizarded tests) where intention would have been visible
AI merges teammates’ contributions	P2	any human witness to the whole
AI backfills absent skills	P12	the team’s knowledge of who could not contribute
The group’s AI dialogue lives in one member’s private account	P15	teammates’ access to their own project’s history
Generated volume buries the human share	P17	which parts a person authored
Fixing and tuning AI output leaves no trace of its own	P17	the judgment-bearing labor behind the polished output
AI may reproduce existing work unnoticed	P14, P15	the ability to check the claim to novelty

on paper before code. With building this easy, teams skip to the thing itself, “*that kind of led us down the path of, like, hey, if we can do it, might as well just do it*” (P5, *on building the real interaction rather than faking it first*). P11’s interface went from “*three rectangles on a page*” (P11) to generated screens with nothing between. The artifact arrives first, and the record of intention forming never exists.

At the execution end, the check that would tie the artifact back to an intention is sometimes absent, “*none of this code is verified... the entire code for this project was pretty much AI-generated*” (P11), and “*I only care for the outcome. I do not know the outcome before I run the code*” (P8). Deliverables accumulate while direction does not. The student stays busy while the architecture supplies and shapes the intentional core of the work. This is the dot’s condition, in which the shaping forces are not felt from inside and the work feels like one’s own.

2.4 Finding 3: The group record dissolves

In team projects the erosion compounds, because a model, or several, stands between the members. One participant described AI as the integrator of the group’s work, the place where teammates’ contributions became one artifact:

“I think I cannot connect the functions together without AI. Coding would not have been possible — we always used AI to merge teammates’ features.” — P2

The model also substitutes for members: “*if my teammates can’t do that, maybe I will think, oh, that’s great, I don’t have to do the design part — but if I have to do the design part, I will ask AI to help me*” (P12). Whoever cannot contribute is backfilled by the tool, invisibly to the rest. The dialogue with the model can itself be private to one member. A flow chart P15’s team used was made in a teammate’s account, in a tool P15 does not use: “*this was on [teammate]’s Gemini. I’m not a Gemini user at all*” (P15). Part of the group’s process history is kept where the other members cannot

reach it, and whether a contribution was made with intention or just quickly generated often goes unsaid.

Asked to divide credit, participants produce splits that do not add up. P11 offered two in the same breath, “*subjectively, split it 50-50, me and my group partner*” (P11) against an objective split in which the model takes the larger share, and noted that the two did not agree. P8, asked about a team idea, said “*the whole idea is not from me... most of the ideas are from myself, when it comes to teamwork, but this one is not*” (P8). One split did add up:

“I think it’ll honestly be 50-50. But, like, 50 between the teammates, the four of us, and then 50 for the AI — so the AI would technically get more per person, I guess.” — P16

In this four-person team the model out-earns every human member. The ghostwriter effect documented for individual writing [6] here takes a collective form, a “we” that cannot enumerate itself. Sometimes the “we” includes work that was never anyone’s. One participant realized that a project they had thought unique reused solutions already online, and called it close to plagiarism (P14, from interview notes).

In other cases students felt that the model’s ideas, documentation, or implementation were better than their own, and they turned to *design intent cosmetics*, work on the signifiers of design work that restores the feeling of authorship. A team redrew AI-generated storyboards by hand:

“It made us feel like we didn’t really do the work, even though this was all our idea, and we just asked it to generate it. At least we’re drawing it up — we’re putting in some effort.” — P15

The process record is also made after the fact. P11’s repository presents a first prototype sketch that was generated by AI after the device was finished, from a photograph of it: “*this picture was made*

after we had already kind of built everything" (P11). The documented process runs backwards, from artifact to sketch. Even inside the record, the source can be lost: "I think this picture is generated, but I'm not sure. Maybe it's from my teammates" (P12), said of an image in the team's own documentation. Another participant scrubbed generated text of its tells, "get rid of all the M dashes... I purposely make it sound like my own writing" (P11). Recent workplace studies report the same practice among professionals, where concealing AI use well is itself treated as a form of AI literacy [28]. A third deleted chat logs by habit, "I just do that, I don't know why, it's so stupid" (P5). None of this changes what the model contributed. It changes how the contribution feels and what remains legible of it to teammates, to graders, and to the students' own later selves. Authorship here is re-enacted over work whose history is already gone. One participant stated the mechanism in a few words:

"the more I work with AI, the less I feel ownership. I know what I want to do, ChatGPT knows exactly how to implement it. If you give LLM control of what to implement, you lose ownership" — P14, from interview notes

Table 1 collects these practices. None of them was concealed from us, and students described each one openly, with ordinary motives. Design intent cosmetics can be individually rational and collectively destructive, since each act restores one person's relation to the work while removing what the group would need to reconstruct it. The cosmetics restore the feeling of authorship and not the record, and what they produce is *design authorship patheticism*, authorship asserted over a history the author can no longer see. CSCW has long shown that cooperative work runs on labor that stays invisible [25], and here the invisible labor includes the acts that decide what the visible record shows. Appeals to honesty will not restore a record that honest people are removing, replacing, or never creating. For the workshop's question, who did what and who is accountable, these cases mark a boundary condition. Attribution presupposes a record, and in AI-saturated group design work that record is dissolving.

ACKNOWLEDGMENTS

We want to thank all the interview study participants from the IDD class in Fall 2025 for their valuable contributions.

REFERENCES

- [1] ACM Publications Board. 2023. ACM Policy on Authorship. <https://www.acm.org/publications/policies/new-acm-policy-on-authorship>. Accessed: 2024-9-5.
- [2] Yoana Ahmetoglu, Marios Constantinides, and Anna Louise Cox. 2026. AI Disclosure with DAISY. In *Proceedings of the 5th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '26)*. ACM, Linz, Austria, 1–23. <https://doi.org/10.1145/3808045.3808056>
- [3] Günther Anders. 1956. *Die Antiquiertheit des Menschen. Band 1: Über die Seele im Zeitalter der zweiten industriellen Revolution*. C.H. Beck, Munich, Germany.
- [4] Lisanne Bainbridge. 1983. Ironies of automation. *Automatica* 19, 6 (1983), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- [5] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX Design Innovation: Challenges for Working with Machine Learning as a Design Material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 278–288. <https://doi.org/10.1145/3025453.3025739>
- [6] Fiona Draxler, Anna Werner, Florian Lehmann, Matthias Hoppe, Albrecht Schmidt, Daniel Buschek, and Robin Welsch. 2024. The AI Ghostwriter Effect: When Users Do Not Perceive Ownership of AI-Generated Text But Self-Declare as Authors. *ACM Transactions on Computer-Human Interaction* 31, 2 (2024), 1–40. <https://doi.org/10.1145/3637875>
- [7] K. J. Kevin Feng, Maxwell James Coppock, and David W. McDonald. 2023. How Do UX Practitioners Communicate AI as a Design Material? Artifacts, Conceptions, and Propositions. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. ACM, Pittsburgh PA USA, 2263–2280. <https://doi.org/10.1145/3563657.3596101>
- [8] Michael Gerlich. 2025. AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking. *Societies* 15 (2025). <https://doi.org/10.3390/soc15010006>
- [9] Yunning Han, Ziyi Qiu, Jiale Cheng, and Ray Lc. 2024. When Teams Embrace AI: Human Collaboration Strategies in Generative Prompting in a Creative Design Task. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–14. <https://doi.org/10.1145/3613904.3642133>
- [10] Emma Harvey, Allison Koenecke, and Rene F. Kizilcec. 2025. Don't Forget the Teachers: Towards an Educator-Centered Understanding of Harms from Large Language Models in Education. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM. <https://doi.org/10.1145/3706598.3713210>
- [11] 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. <https://doi.org/10.48550/arXiv.2607.26387> [cs.CY] Workshop at the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW '26).
- [12] Ahmed Kharrufa and Ian Johnson. 2024. The Potential and Implications of Generative AI on HCI Education. In *Proceedings of the 6th Annual Symposium on HCI Education*. ACM, New York NY USA, 1–8. <https://doi.org/10.1145/3658619.3658627>
- [13] Hao-Ping Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, 1–16. <https://doi.org/10.1145/3706598.3713778>
- [14] Lawrence Lessig. 2006. *Code: Version 2.0*. Basic Books, New York, NY, USA.
- [15] Jie Li, Hancheng Cao, Laura Lin, Youyang Hou, Ruihao Zhu, and Abdallah El Ali. 2024. User Experience Design Professionals' Perceptions of Generative Artificial Intelligence. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–18. <https://doi.org/10.1145/3613904.3642114>
- [16] Jie Li, Youyang Hou, Laura Lin, Ruihao Zhu, Hancheng Cao, and Abdallah El Ali. 2026. Vibe Coding in Product Teams: Reconfiguring AI-Assisted Workflows, Prototyping, and Collaboration. In *Proceedings of the 5th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '26)*. ACM, Linz, Austria, 1–16. <https://doi.org/10.1145/3808045.3808062>
- [17] Lujin Mao, Mengyao Qi, and Zhibin Zhou. 2026. A Review of Generative AI Integration in Design Education: Macro and Micro Perspectives on University Policies and Course Practices. In *Proceedings of the 2026 Designing Interactive Systems Conference*. <https://doi.org/10.1145/3800645.3813017>
- [18] Karl Marx and Friedrich Engels. 1848. *Manifest der Kommunistischen Partei*. Office der Bildungs-Gesellschaft für Arbeiter, London, UK. English: The Communist Manifesto.
- [19] Maria Teresa Parreira, Pranav Prabhat Sinha, Hauke Sandhaus, and Wendy Ju. 2026. Visible Adoption, Hidden Contribution: Three Cohorts of AI Use in an HCI Prototyping Course. In *Symposium on the Future of Computer Work (CHIWORK '26)*. Linz, Austria. <https://arxiv.org/abs/2606.14054> Workshop paper.
- [20] Paul C. Parsons. 2026. Myths and Ironies of AI-Assisted Design. In *Proceedings of the 2026 Designing Interactive Systems Conference*. <https://doi.org/10.1145/3800645.3813086>
- [21] Jonathan Pugh. 2020. *Autonomy, Rationality, and Contemporary Bioethics*. Oxford University Press, Oxford, UK. <https://doi.org/10.1093/oso/9780198858584.001.0001>
- [22] Hauke Sandhaus, Qiuquan Gu, Maria Teresa Parreira, and Wendy Ju. 2025. Co-Designing with Transformers: Unpacking the Complex Role of GenAI in Interactive System Design Education. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 1228–1243. <https://doi.org/10.1145/3715336.3735805>
- [23] Hauke Sandhaus, Kashif Imteyaz, Mohammed Almutairi, Pooja Prajod, Divya Ramesh, Saiph Savage, Qian Yang, and Michael Muller. 2026. Interrogating GenAI Augmentation for CHIworkers: Strategies for Professional Autonomy and Accountability. In *CHIWORK '26: Proceedings of the 5th Annual Symposium on Human-Computer Interaction for Work (Workshop)* (Linz, Austria) (CHIWORK '26). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3805029.3818271>
- [24] Joon Gi Shin, Janin Koch, Andrés Lucero, Peter Dalsgaard, and Wendy E. Mackay. 2023. Integrating AI in Human-Human Collaborative Ideation. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, 1–5. <https://doi.org/10.1145/3544549.3573802>

- [25] Susan Leigh Star and Anselm Strauss. 1999. Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work. *Computer Supported Cooperative Work (CSCW)* 8, 1–2 (1999), 9–30. <https://doi.org/10.1023/A:1008651105359>
- [26] Dakuo Wang, Justin D. Weisz, Michael Muller, Parikshit Ram, Werner Geyer, Casey Dugan, Yla Tausczik, Horst Samulowitz, and Alexander Gray. 2019. Human-AI Collaboration in Data Science: Exploring Data Scientists' Perceptions of Automated AI. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 211:1–211:24. <https://doi.org/10.1145/3359313>
- [27] Langdon Winner. 1977. *Autonomous Technology: Technics-out-of-Control as a Theme in Political Thought*. MIT Press, Cambridge, MA, USA.
- [28] Qing Nancy Xia, Marios Constantinides, Advait Sarkar, Duncan P. Brumby, and Anna L. Cox. 2026. "If You're Very Clever, No One Knows You've Used It": The Social Dynamics of Developing Generative AI Literacy in the Workplace. In *Proceedings of the 5th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '26)*. ACM, Linz, Austria, 1–18. <https://doi.org/10.1145/3808045>
- [29] Nur Yildirim, Alex Kass, Teresa Tung, Connor Upton, Donnacha Costello, Gene Green, Tingting Dong, and John Zimmerman. 2022. How Experienced Designers of Enterprise Applications Engage AI as a Design Material. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA. <https://doi.org/10.1145/3491102.3517491>
- [30] J D Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: How non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–21. <https://doi.org/10.1145/3544548.3581388>
- [31] Shuning Zhang, Hui Wang, and Xin Yi. 2025. Exploring Collaboration Patterns and Strategies in Human-AI Co-creation through the Lens of Agency: A Scoping Review of the Top-tier HCI Literature. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (2025), 1–43. <https://doi.org/10.1145/3757594>