Supplementary Chapters: An Appendix for the Dissertation

by

Larissa Schwartz

B.A., University of Iowa, 2007

M.S., University of Colorado Boulder, 2023

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Engineering Education

2026

Committee Members:

Tom Yeh, Chair

Diane Sieber

Lee Frankel-Goldwater

Bo Wu

Angela R. Bielefeldt

Schwartz, Larissa (Ph.D., Engineering Education)

Supplementary Chapters: An Appendix for the Dissertation

Thesis directed by Prof. Tom Yeh

This appendix contains five supplementary chapters that complement the main dissertation document, together with an evidence table catalogue and an AI use disclosure. The main document operates under design-based research (McKenney and Reeves 2018) grounded in constructivist learning theory and self-determination theory, and surfaces three curriculum-design principles: modularity, learner choice, continuous feedback. This appendix supplements that primary methodology with an analytic-autoethnographic posture (Anderson 2006), introduced in Appendix B and carried through Appendices C and D. The autoethnographic supplement does not introduce competing findings; it elaborates each of the three principles by drawing on the practitioner-pioneer's reflexive position, naming the patterns within each principle that the DBR analysis does not surface: compression-as-curriculum-maturation (the elaboration of modularity), multi-channel teaching practice (the elaboration of learner choice), and hallucination-as-pedagogy (the elaboration of continuous feedback). Appendix A establishes practitioner-pioneer positionality and the four-theme curriculum architecture; Appendix E is the master evidence catalogue; Appendix F discloses AI use during dissertation preparation.

Note

This document contains the appendix material of the dissertation. It is intended to be read alongside the main dissertation document.

Preliminary note on identifiability and PII review. This appendix is being circulated to the committee as a draft. It contains material that has not yet completed a final review for personally identifiable information. Before the appendix is finalized for submission, every mention of a named individual will be reviewed against the taxonomy in §B.6.5: students named in instructor-produced materials will be anonymized unless explicit written consent for educational use is documented; named guest speakers will be retained as public professional identities with their professional context attached; Luma-platform workshop feedback will be reviewed for anonymization; and external journalism is retained as already published and consented.

Acknowledgements

Acknowledgements appear in the main dissertation document.

Contents

Chapter

Appendix

Appendix A

Positionality and Curriculum

A.1 Literature survey

In this chapter I situate myself as a practitioner-pioneer of generative-AI instruction at a research university, and I situate the four-theme curriculum architecture I built as the analytic lens this appendix uses to supplement the main document's curriculum-design principles. The literature survey is short because the appendix's supplementary posture rests on a single framework (analytic autoethnography per Anderson 2006, established in Chapter B alongside the main document's design-based research). The survey's purpose is to position the pioneer instructor as an autoethnographic site worth studying and to locate the gap that the appendix fills.

A.1.1 Analytic autoethnography as the supplementary framework

The main document of the dissertation operates under design-based research (McKenney and Reeves 2018) grounded in constructivist learning theory and self-determination theory. This appendix supplements that primary methodology with analytic autoethnography (Anderson 2006). Chapter B lays out the supplementation rationale and how the two methodologies operate together. The relevant point for this chapter is that analytic autoethnography is the framework that lets me treat my own practitioner-pioneer position as an analytic resource alongside the DBR-grounded curriculum-design principles, surfacing theoretical findings the DBR analysis developed in the main document does not surface. The chapters that follow apply the autoethnographic supplement to my biography and curriculum (this chapter), to methodology (Chapter B), to the four iterations

(Chapter C), and to the theoretical findings the autoethnographic analysis surfaces (Chapter D).

A.1.2 The pioneer instructor as an autoethnographic site

The literature on practitioner-pioneer instruction in emerging-technology fields is thin. Generative-AI pedagogy at the university level is a young area; the first systematic course offerings emerged in 2023, and the scholarly literature on how those first courses were taught is mostly forthcoming. Cochran-Smith and Lytle (2009) argue more broadly that practitioner inquiry is an underused source of theoretical insight in education research, and that the practitioner's own position carries information that external research cannot recover. Anderson's (2006) framing of the complete member researcher is a methodologically specific instance of the same claim.

I draw on this background to make a simple positioning claim. My pioneer-instructor role is itself a site at which analytic autoethnography can do work that other methodologies could not. The four iterations of my course are not a controlled experiment, and they are not a comparative study of multiple instructors. They are a single practitioner's pioneer practice, documented from the inside, across three years and multiple institutional contexts. The literature has few examples of this configuration treated analytically.

A.1.3 Generative-AI pedagogy in higher education

The early literature on generative-AI in higher education has moved quickly. The initial wave of scholarly response (largely from 2023 and 2024) focused on academic-integrity questions, on the system limitations of large language models, and on policy and governance frameworks for institutional adoption. A second wave attended to curriculum design and instructional practice: Mollick and Mollick (2023) offered the most widely-circulated practitioner-facing essay on assigning AI to undergraduate work; Long and Magerko's (2020) earlier "What is AI literacy?" framework was retrofitted by many instructors to the generative-AI moment; Touretzky, Gardner-McCune, Martin, and Seehorn's (2019) AI4K12 "five big ideas" framework remained the dominant K-12-facing reference. Most of the published work in this second wave is conceptual or framework-oriented

rather than empirical.

The empirical literature on what actually happened when an instructor designed and delivered a generative-AI course from scratch in 2023 or 2024, on how the course evolved across iterations under tool turnover, on how learners across age groups engaged with the same conceptual architecture, and on how the instructor's practice spread across multiple delivery channels, is what my dissertation contributes to. I claim a documented practitioner-pioneer record (Iterations 1 through 4, plus six additional delivery channels) of a kind that the early literature has not yet produced.

A.1.4 The gap I locate

I locate the gap as follows. The methodology literature offers analytic autoethnography (Anderson 2006) as a framework for treating the complete-member researcher's insider position as analytic data. The practitioner-inquiry literature (Cochran-Smith and Lytle 2009) argues that practitioner positions hold information unavailable to external research. The generative-AI pedagogy literature, however, has not yet produced empirical practitioner-pioneer records of the multi-iteration multi-channel kind. My dissertation occupies that gap. It applies an established methodology to a new site (generative-AI pedagogy at a research university in 2023 through 2025), documents the pioneer-instructor practice across four iterations and eight delivery channels, and develops three nameable theoretical findings that emerge from the cross-iteration analysis.

That is the contribution claim my dissertation makes. Subsequent sections of this chapter develop the practitioner-pioneer biography (§A.2), the institutional positioning (§A.3), the four-theme curriculum architecture (§A.4), and the cross-context theme stability (§A.5) that together establish the work as standalone scholarship within engineering education and HCI.

A.2 Practitioner-pioneer biography

In this section I narrate the path that brought me to the work the dissertation documents. The narration is drawn primarily from my Research Impact Essay (RE), which I wrote in the first person near the end of Iteration 1 and which carries my own account of how I arrived at the CU

course. The biography matters analytically because the framework I built (Education, Industry, Ethics, Accessibility) and the way I taught it carry the imprint of where I came from.

A.2.1 My background as ELA and high-school art teacher

I trained and taught as an English-language-arts and high-school art teacher before entering the doctoral program at CU Boulder. My classroom experience was secondary-level work in subject matter that prized open-ended creative production, individual student voice, and the rhetorical and ethical dimensions of made objects. The eye I brought to generative-AI tools when they emerged was a teacher's eye on what those tools made possible for students, not a researcher's eye on what those tools failed at as systems.

A.2.2 Spring 2023 · the Colorado school-district professional development

In the spring of 2023 I led a professional development session for a school district in Colorado on how teachers and students could use generative-AI art applications, including DALL-E, Midjourney, and NightCafe, to create artwork from prompts. I wrote about this work in my Research Impact Essay:

> "Over a year ago I started researching Generative AI and how teachers and students can use different art applications, such as DALL E, Midjourney and NightCafe in order to create artwork from a prompt." (RE-Q1)

> "After creating an art curriculum with these different AI art generators, I shared my ideas with a school district in Colorado, where I led a professional development to teach the staff how to incorporate and collaborate with these tools in the classroom." (RE-Q2)

The school-district professional development was my first sustained generative-AI teaching practice. It was a teacher-to-teacher exchange focused on how to bring the tools into existing curricula rather than on what the tools could or could not do as systems. The professional-development orientation is visible in the four-theme architecture I subsequently built at CU: each

theme is named in terms of what generative AI does for a domain of practice (Education, Industry, Accessibility) rather than in terms of how the underlying models work.

A.2.3 The Charles Burrell School of Arts diptych contest

Following the professional development I designed and ran an art contest for high-school students at the Charles Burrell School of Arts. Each student produced two pieces of work: a hand-drawn or hand-painted 2-D piece and a generative-AI-prompted piece, displayed together as a diptych. The contest title was Human vs AI, and the diptych format made the comparison concrete and student-facing.

The contest drew more than sixty competitors. In my Research Impact Essay I wrote:

> "From this art contest, we had over 60 students compete and it was such a success that it motivated me to pilot the first ever Generative AI class at the University of Colorado Boulder through the Atlas Institute." (RE-Q3)

This is the autoethnographic finding I name at the opening of the dissertation. The arc was not the conventional researcher-to-teacher arc, in which a graduate student designs an experiment and then teaches a course to gather data. The arc was a teacher-to-researcher arc, in which a teacher's classroom practice in K-12 motivated her to pilot an early generative-AI course at her university. (I use my own first-person language above; "first ever" is my framing at the time of writing the Research Impact Essay. Without an institutional registry of GenAI course offerings nationally, the precise pioneering claim I can support is that this was among the earliest systematic offerings at CU Boulder and in the field nationally during 2023-2024, not that it was first absolutely.) The data of the dissertation is not separable from the biography that produced it.

A.2.4 The CU pilot and the four-theme architecture I built

Following the Charles Burrell contest I designed the curriculum that became Iteration 1 at the ATLAS Institute. I named the curriculum architecture explicitly in my Research Impact Essay:

> "During this spring semester, I've been teaching Generative AI to undergraduate students about how these different applications can be used in Education, Industry and Accessibility, and the ethical concerns that arise while using these programs." (RE-Q4)

This is a tighter rendering of the four themes than the dissertation's earlier framing. The Research Impact Essay names Education, Industry, and Accessibility as the three standalone themes, with Ethics as the cross-cutting fourth theme that runs through each of the other three. The dissertation's earlier four-theme rendering treats Ethics as a standalone fourth theme, and the curriculum as actually taught reflects this earlier framing (the Iteration 2 syllabus SY-2 schedules Weeks 1-4 as Education, Weeks 5-11 as Industry, Weeks 12-15 as Accessibility, and Ethics as cross-cutting across all). The two renderings are compatible, and the cross-cutting status of Ethics is itself a finding worth marking in §A.4 below.

A.2.5 What this biography means analytically

The K-12-to-CU arc has three analytic consequences for the dissertation.

First, it grounds the four-theme architecture in actual pedagogical practice rather than in a literature-derived analytic frame. The themes were how I taught the work from the beginning. They are not a researcher's overlay imposed on a course taught for other reasons.

Second, it predicts the multi-channel pattern that Chapter D develops as a substantive finding (§D.4). A teacher who came to generative-AI work through K-12 professional development and a high-school art contest is a teacher who is likely to keep K-12 outreach alongside her university teaching. The STEAM Festival mural (ST-MURAL), the UW KidsTeam research collaboration (KT-DECK and the KT corpus), and the storytelling-with-Cartoonimator worksheet (STC) are continuous with the K-12 origin rather than discontinuous from it.

Third, it shapes the practitioner-pioneer position I claim. I claim to be among the earliest systematic offerors of a generative-AI class at CU Boulder during the post-ChatGPT period (2023-2024), and among the early systematic offerings in the field nationally. The pioneer claim is not biographical decoration; it is the precondition for the methodological framework Chapter B adopts.

Analytic autoethnography requires complete member researcher status at the setting being studied. My pioneer status at the CU generative-AI teaching site is what licenses the methodological choice, and the claim does not depend on being absolutely first; being among the earliest is sufficient.

The biography is short, sourced primarily to my own first-person essay (RE), and consequential for the chapters that follow.

A.3 Institutional positioning

My work crossed institutional boundaries within CU Boulder and reached beyond CU through partnerships with other universities and external programs. In this section I document where the work was hosted and funded at each stage, because the institutional positioning is itself a finding about how pioneer practice survives and propagates at a research university.

A.3.1 My doctoral home in ENED

My doctoral program is Engineering Education (ENED) within the College of Engineering and Applied Science at CU Boulder. The dissertation cover, the Iteration 2 course code (GEEN 3830-001, documented in SY-2), and the Preliminary Exam Part 2 document (PR-PART2) all locate my program in ENED.

I name this because my dissertation is not a computer-science thesis. It is an engineering-education thesis about a generative-AI course I taught. The disciplinary frame is engineering education and adjacent HCI work; the analytic posture is analytic autoethnography of my own pioneering practice; the contributions are calibrated to engineering-education and HCI literatures.

A.3.2 My lab affiliation at the Imagine AI Lab

My laboratory affiliation is with the Imagine AI Lab, led by Prof. Tom Yeh, my PhD thesis advisor. The ForeverGold deck for Iteration 1 (DK-1.FG) names the affiliation on its title slide. The lab's research interests in human-AI interaction and generative-AI tooling provide the immediate intellectual setting for my doctoral work, and Prof. Yeh's advising directly shapes the dissertation.

The Imagine AI Lab affiliation also connects to the K-12 strands that surfaced across the cross-iteration corpus (ST-MURAL, ST-PHOTO, KT corpus, STC). The lab's interest in how learners across age groups engage with AI is what made my K-12 art-teacher origin (RE-Q2 through RE-Q3) continuous with my doctoral research practice rather than discontinuous from it.

A.3.3 The AI-IRT Seed Grant as foundational context

The AI-IRT Seed Grant (AI-PROPOSAL) provided the institutional and funding context within which my dissertation work unfolded. The grant's principal investigators are Tom Yeh (Computer Science) and Diane Sieber (Herbst Program for Engineering, Ethics, and Society). Both PIs serve on my dissertation committee.

The grant's scope explicitly covered K-12, undergraduate, and graduate strands of AI literacy work. My dissertation occupies the undergraduate strand most directly, but the K-12 strand maps onto my Charles Burrell School of Arts contest (RE-Q3) and my STEAM Festival mural (ST-MURAL), and the graduate strand maps onto my HCI summer 2024 guest-lecture series (HC corpus). The grant funded both me and another doctoral student (Mohsena Ashraf) as PhD researchers.

I name the grant context to make a positioning point that the original dissertation draft underplays. My pioneer work did not happen in isolation. It happened within an institutionally-seeded research program co-led by my committee chair (Tom Yeh) and one of my committee members (Diane Sieber), and the grant's scope explicitly authorized the multi-channel work that Chapter D documents as a substantive finding.

A.3.4 Iteration-specific institutional sponsors

The four iterations of my course were hosted under different institutional banners within CU Boulder and through one external partnership.

Iteration	Period	Institutional sponsor	Course or program designation
Iteration 1	Spring 2024	ATLAS Institute / CTD	CTD pilot course
Iteration 2	Spring 2025	College of Engineering (CEAS)	GEEN 3830-001 Special Topics
Iteration 3	August 2025	College of Engineering and Applied Science (CEAS)	GenAI in Five online workshop, first cohort
Iteration 4	September 2025	GenAI Works (external partnership)	GenAI Works cohort, YouTube-delivered

The cross-iteration institutional movement is a finding in its own right. My course began at ATLAS as an experimental CTD pilot, moved into the College of Engineering as a credit-bearing special-topics course, compressed into a CEAS-sponsored online workshop, and was finally delivered through a non-CU partnership that reached a global audience. Each move expanded the institutional reach of the curriculum while the curriculum architecture itself remained stable (§A.4 below; §C.6.5 elsewhere).

A.3.5 External and partner institutions

Beyond CU Boulder, my practice crossed into other institutional contexts.

University of Washington · KidsTeam and YAB. In July 2024 I participated as a CU-based researcher and collaborator in three-day KidsTeam co-design sessions and multiple Youth Advisory Board (YAB) sessions at the University of Washington (KT corpus). The collaboration produced K-12 research data on children's and teens' reasoning about generative AI in schools (KT-THEMES).

GenAI Works. Iteration 4 (September 2025) was hosted via partnership with GenAI Works, which delivered the workshop through its YouTube channel to a global audience (TR-4.D1 through TR-4.D5). The Day 1 transcript names attendees joining from Nigeria, the United Kingdom, Denver, and Costa Rica within the first thirty lines of session opening.

Rocky Mountain Advanced Computing Consortium and NAIRR Pilot. In March 2026 I delivered a webinar titled "Unleashing Creativity with Generative AI" for the Rocky Mountain Advanced Computing Consortium (RMACC). The webinar was uploaded by the federal NAIRR Pilot program (National AI Research Resource), bringing my work to a federal research audience (WB-2026-03-03).

A.3.6 Boundary-crossing as positionality

The boundary-crossing pattern is itself a positionality finding. I am a doctoral student in engineering education (ENED), lab-affiliated with an arts-and-creativity institute (ATLAS), funded under a seed grant co-led by computer science and engineering ethics (AI-PROPOSAL), teaching first under CTD and then under engineering course numbers (GEEN 3830-001), with external partnerships at a peer research university (UW KidsTeam) and a private education platform (GenAI Works), and delivery audiences ranging from K-12 children at the STEAM Festival to federal research audiences via the NAIRR Pilot.

For my positionality statement (folded into §B.3 of the main document), this means I do not occupy a single disciplinary or institutional position. I occupy a network of positions that sit at the intersections among engineering education, HCI, arts-and-creativity research, and K-12 outreach. The work the dissertation documents is what that network produced.

A.4 The four-theme curriculum architecture

In this section I define the four-theme curriculum architecture that I built into my course from Iteration 1 onward. The architecture is the analytic spine of the dissertation. It is not a researcher's overlay on a course taught for other reasons; it is the framework I taught from the beginning.

A.4.1 The four themes

The four themes are Education, Industry, Ethics, and Accessibility. Each is defined below in terms of what generative AI does for a domain of practice, which is the framing I used in my own first-person essay (RE-Q4) and which appears across my Iteration 1 ForeverGold deck (DK-1.FG, slides 5-9), my Iteration 2 syllabus (SY-2), and my Iteration 1 Final Project Requirements (FP-1).

Education

Education names the use of generative-AI tools within instructional settings: by teachers in lesson planning and material generation, by students in coursework and study, and by educational designers in curriculum design. The theme covers prompt engineering as a teachable skill, tool integration into existing curricula (drawing on my Charles Burrell professional-development experience, RE-Q2), and the question of how to teach with rather than against generative-AI tools. The Education theme is the opening territory in every iteration. Week 1 of Iterations 1 and 2 (DK-1.W01 and DK-2.JAN13) is the prompt-engineering opener; the early weeks of both semester iterations cover Education before moving outward into the other themes.

Concrete example. Iteration 1 Week 1 (January 17, 2024) opened with a prompt-engineering lecture and an interactive assignment in which students wrote five facts about themselves, used ChatGPT to generate multiple-choice quizzes from those facts, used DALL-E 3 via Microsoft Designer to generate an "alter-ego animal" image of themselves, and exchanged the quiz and image with a partner as an introduction activity (per the Week 1 entry in WU-1.W01). The activity instantiates the Education theme on three levels at once: it is a teacher's lesson plan using a generative-AI tool, it is a student-facing assignment in prompt engineering, and it is curriculum design with the tool as a structural element rather than as a peripheral exercise. Twenty-three named student outputs from the session are preserved in DK-1.W01.

Industry

Industry names the use of generative-AI tools in professional work outside the classroom. The

theme covers software engineering (DK-2 has Justin Shacklette as the GenAI-for-Software guest in Iteration 2), data tooling and notebook environments (Bobby Hodgkinson on NotebookLM in Iteration 2), and the practical workplace applications students would encounter as they entered industry. Iteration 1 brought industry voices through Daniel Ritchie's Hugging Face workshop and Matt Zago's video work; Iteration 2 brought industry voices through Tom Yeh on DeepSeek and Justin Shacklette on software engineering. The Industry theme operates as the practical-application bridge between Education and the more critical themes that follow.

Concrete example. Iteration 1 Week 5 was the opening week of the Industry theme block. I delivered a lecture on how generative AI is used across professional industries and then split the students into small groups, with each group assigned a specific industry application category to research and report back to the class. The categories named in WU-1.W05 include Augment Data, Synthetic Data, Drug Design, Design Neural Network, Chip Design, Create Algorithm, Design of Parts, 3D Shape Creation, Create Text, Increase Image Resolution, Creation of an Instance Image, Image-to-Image Conversion, Text-to-Speech Generator, Create Music, Generate Videos, Generate Image, and Material Science. The activity converts the Industry theme from a topical block into a student-driven survey of generative-AI applications across professional domains, with each student becoming the class's local expert on one application category.

The curated industry network as part of the contribution. Beyond the activity structure, my Industry-theme contribution includes the network of practitioners I curated to bring into the classroom across the two semester iterations. Daniel Ritchie (Hugging Face workshop, Iteration 1) brought open-model practice into reach of undergraduates; Matt Zago (Iteration 1) brought finance-industry generative-AI applications; Nikolaus Klassen (Google, Iteration 1) brought industry-internal ethics; Justin Shacklette (Iteration 2) brought GenAI for software engineering; Bobby Hodgkinson (Iteration 2) brought NotebookLM and the data-tooling perspective. The roster was not a fixed lecture series. It was an actively rotating network I assembled, replenished, and re-aimed at each iteration's evolving technology landscape. The curation work itself is part of what the Industry-theme contribution consists of: students did not encounter generative AI as an abstract

category but as an industry-applied practice represented by named practitioners I had identified, invited, and embedded into the course. The cross-iteration guest-speaker turnover documented in §C.6.3 makes the curation visible as an instructional practice in its own right.

Ethics

Ethics names the cross-cutting attention to questions of consent, harm, fairness, intellectual property, environmental cost, and human creativity that generative-AI tools raise. In my Research Impact Essay I named Ethics as "the ethical concerns that arise while using these programs" (RE-Q4), framing it as a cross-cutting theme that runs through Education, Industry, and Accessibility rather than as a standalone theme of equal kind. The Iteration 2 syllabus (SY-2) is consistent with this framing: it schedules Education, Industry, and Accessibility as discrete blocks while marking Ethics as a thread that runs through all of them. The Iteration 1 Final Project Requirements (FP-1) operationalize this with a dedicated ethics question, "Do you think this will eliminate creative jobs?", that runs alongside the technical and applied questions in the same assignment. I treat Ethics as the architecture's cross-cutting fourth theme in this dissertation, while acknowledging that the curriculum-as-taught sometimes treats it as a standalone block (Iteration 1 Weeks 10-12 with Klassen Google ethics and Weng Reality Editor are the visible exception).

Concrete example. Iteration 1 Week 7 surfaced Ethics within the Industry theme through a deepfakes-and-music sequence. I delivered a lecture on the ethics of generative-AI video and audio (deepfake videos, face swaps, audio deepfakes), discussed the "Heart on My Sleeve" case study in which Drake's and the Weeknd's voices were generatively reproduced without consent, and assigned the students to find a song that had been AI-generated and identify the signals by which they had known it was generated (per WU-1.W07). Students then created their own deepfake videos and still images using ROOP or a generative-AI app of their choice, framed explicitly as ethical deepfakes for educational use, and reflected on the experience in writing. The activity makes Ethics concrete by asking students to enact the technology, surface their own reasoning about its limits, and produce a reflection. The cross-cutting framing means the Ethics work appears inside the Industry theme

rather than in a separate Ethics block.

Accessibility

Accessibility names the use of generative-AI tools in the lives of people with disabilities and in service of equitable access. The theme covers tools designed for accessibility (Be My AI for blind and visually-impaired users, named in SY-2; ElevenLabs and Speechify for audio accessibility), and it brings lived-experience speakers and Disability-Studies framings into the course. The Iteration 1 Research Impact Essay names Accessibility as the third of the standalone themes (RE-Q4), and the Iteration 1 curriculum closes with Accessibility before students undertake the Final Project (FP-1). Iteration 2 schedules Accessibility for Weeks 12-15 (per SY-2). The Accessibility theme also extends naturally into the cross-iteration K-12 work, where children at the STEAM Festival (ST-MURAL) and at UW KidsTeam (KT-DECK, KT-IDEAS) engaged with generative-AI tools in ways that surfaced accessibility considerations.

Concrete example. Iteration 1 Weeks 14 and 15 brought Derek Riemer, a Googler who works on Google Drive's web interface and is blind from birth, into the course as the closing guest before the Final Project. Riemer demonstrated the AI features in Be My Eyes (the precursor product to Be My AI) and walked the students through assistive technologies he uses daily, including screen readers, canes, and AI applications that identify objects from camera input (per WU-1.W14). The students had also, earlier in the Accessibility block (Week 13), practiced descriptive prompting by handwriting a paragraph describing an image, asking a partner to draw the image from the description, and then asking Microsoft Copilot to describe the same image as a human-versus-AI accessibility-of-description exercise. The activity sequence makes Accessibility concrete by combining lived-experience instruction (Riemer's demonstrations) with student practice in the descriptive skills that accessibility work requires.

A.4.2 Where the architecture is explicit in my artifacts

The four-theme architecture is not implicit. It appears explicitly in the artifacts I produced from Iteration 1 onward.

Artifact	Where the architecture appears
ForeverGold deck (DK-1.FG)	Slides 5-9 explicitly name the four themes as the framing for Iteration 1
Iteration 1 Final Project (FP-1)	Nine reflection questions distributed across the four themes; ethics question on creative-jobs displacement appears alongside the other three theme questions
Research Impact Essay (RE-Q4)	"Education, Industry and Accessibility, and the ethical concerns that arise"
Iteration 2 syllabus (SY-2)	Weekly structure: Weeks 1-4 Education, 5-11 Industry, 12-15 Accessibility, Ethics cross-cutting
Iteration 2 slide decks (DK-2 series)	Dated lectures named by theme: Industry theme (DK-2.MAR05), Education theme (DK-2.MAR10)
HCI summer 2024 deck series (HC corpus)	Decks named explicitly by theme: HC-EDU, HC-INDUSTRY, HC-ACCESS, HC-AUDIO, HC-VIDEO, HC-MUSIC plus their ethics-paired counterparts
Keep Up Newsletter (KN-EP1 through KN-EP3)	Topics organized by theme cluster (image, research, sound) within the broader theme architecture
CU RMACC webinar (WB-2026-03-03)	Public-facing synthesis of my work organized around the four themes

The first row of the table above points at slides 5 through 9 of the ForeverGold deck; Figure A.1 shows those slides directly, alongside the title slide for context.

(a) Slide 1 (title)

(b) Slide 5

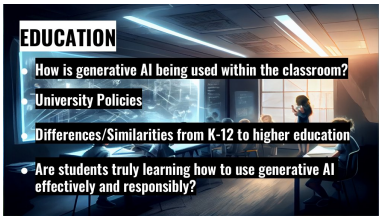

(c) Slide 6

(d) Slide 7

(e) Slide 8

(f) Slide 9

Figure A.1: Selected slides from the ForeverGold course deck (DK-1.FG): the title and the four-theme architecture on slides 5 through 9.

A.4.3 The architecture as analytic lens

For this dissertation, the four-theme architecture is the analytic lens I use to organize the iteration narratives (Chapter C), to identify the cross-iteration patterns (§C.6), and to develop the substantive findings (Chapter D). The framework's stability across the iterations is itself a finding I develop in §A.5 below; the framework's appearance across delivery channels is itself a finding I develop in §D.4.

Importantly, the framework is not an analytic imposition that came after the data. It is the framework I taught from. The dissertation's use of the framework as analytic lens follows from the practitioner-pioneer position I claim: the lens is the lens I used as a practitioner, surfaced now as an analytic instrument because it is the one the data was generated through.

This is the move analytic autoethnography permits and asks for. My insider position is the source of the lens. The lens is then tested against cross-context evidence (children, undergraduates, journalists, online learners), and the testing constitutes the trustworthiness move that §B.6 develops. Section A.5 next presents the cross-context testing for the four-theme architecture.

A.4.4 What is and is not novel about the architecture

The four themes (Education, Industry, Ethics, Accessibility) are not novel as a list. Adjacent frameworks in AI literacy use overlapping categories: Long and Magerko (2020) name 16 AI literacy competencies that include ethics, fairness, and accessibility-relevant strands; Touretzky, Gardner-McCune, Martin, and Seehorn's (2019) AI4K12 framework organizes K-12 AI literacy around five "Big Ideas" that include perception, representation, learning, natural interaction, and societal impact; and disability-studies-informed AI work treats accessibility as a first-order concern. A reader who counts the four themes against these adjacent frameworks will find substantial overlap.

What I claim as novel is not the naming of these four themes but the empirical documentation of how the four themes function **as a stable architecture under tool turnover and delivery-format compression** across four iterations and eight contexts over two and a half

years. The contribution is the record of architectural stability, not the architecture's category labels. A subsequent instructor who adopts the four themes inherits a framework whose presence has been documented across tool generations (DALL-E to Midjourney to Nano Banana; ChatGPT to DeepSeek to Claude 3.7) and audience shifts (CTD undergraduates to mixed-engineering undergraduates to global online learners to children at the STEAM Festival). That stability record is what the dissertation contributes; the theme names are the convenient handles by which the record is indexed.

A.5 Theme stability across contexts

In this section I test the four-theme architecture against learners and observers outside my own classroom. The test matters because if the architecture only appears in materials I produced, the architecture could be read as a researcher's overlay; if the architecture appears in independent voices across age groups and contexts, the architecture has external validity.

I present five sources of cross-context evidence: K-12 children, university undergraduates, journalists, online workshop attendees, and my own public-facing writing across delivery channels.

A.5.1 K-12 children at UW KidsTeam surface the same themes in a co-design context

In July 2024 I participated in three days of UW KidsTeam co-design sessions with children and multiple Youth Advisory Board (YAB) sessions with teens at the University of Washington. The sessions investigated children's and teens' opportunities and challenges with generative AI in schools. The KidsTeam methodology was led by the UW team; I was a participating collaborator, not the facilitator. I did not bring my four-theme framework into the sessions as a category for the children to use; the framework's appearance in their reasoning is therefore **near-independent** of my own curriculum — the methodology was someone else's, the children were not my students, and the theme analysis (KT-THEMES) was the UW team's. What it is not is fully independent observation, since I was in the room. With that caveat in place, the convergence between the children's themes and the architecture I had built is meaningful evidence for the architecture's

reach beyond my classroom.

The KidsTeam research themes document (KT-THEMES) catalogues what children and teens surfaced. Within "challenges," they named hallucinations such as "images produced with a third arm" (KT-THEMES-C5), which maps to the Ethics theme as a question about reliability and to the Accessibility theme as a question about whether AI-generated images can represent human bodies fairly. They named lack of emotion and human connection (KT-THEMES-C4), which maps onto the Ethics and Education theme cluster. They named cheating (KT-THEMES-C2), which is a recurring Ethics concern in classroom contexts.

Within "opportunities," they named teachers allowing students to use generative AI for paragraphs with mistakes to correct (KT-THEMES-O1), which maps onto the Education theme. They named producing faster work (KT-THEMES-O3), which maps onto the Industry theme through the productivity frame. They named checking math homework (KT-THEMES-O2), which maps onto the Education theme as an instructional application.

The children and teens at UW were not coached toward my framework, and the framework's themes appear in their reasoning. This is the strongest external-validation evidence the cross-iteration corpus carries.

A.5.2 Undergraduate cohorts in my own courses apply the framework in final-project work

The Iteration 1 Final Project Requirements (FP-1) distribute nine reflection questions across the four themes. Students who completed the final project were therefore prompted to engage with all four themes, and their work is part of the dialogue-with-informants-beyond-self evidence base.

The Iteration 2 student teach-out presentations (SP-2 series) extend this by giving students an opportunity to teach back. The recorded teach-outs cover DeepfakeAI (SP-2.DEEPFAKE, an Ethics topic with Industry implications), Sintra.ai (SP-2.SINTRA, an Industry topic), Haley Phillips's teach-out (SP-2.HALEY), Dakota A's teach-out (SP-2.DAKOTA), and AI and Robotics integration (SP-2.ROBOTICS, an Industry topic with Education implications). The distribution

of student-chosen topics across the four themes is a piece of evidence that the framework matched what students cared about, rather than something I had to insist on.

Specific named undergraduates also surface in the cross-corpus material in ways that document theme engagement. Ashley Stafford appeared in the Iteration 1 Week 1 deck (DK-1.W01) and was subsequently quoted in the Aspen Public Radio article (AP-2024-05-16) as the named student voice on what the course offered. Ethan Cuenca appeared in the Iteration 1 Week 1 deck (DK-1.W01) and was publicly credited in my Keep Up Newsletter Episode 3 (KN-EP3-Q1, "I learned about Soundful from one of my students during class") as the student whose tool discovery shaped Iteration 2's curriculum. Both students are documented engagements with the framework across multiple artifacts.

A.5.3 Journalists frame my work in theme terms

The Aspen Public Radio article (AP-2024-05-16), published one week after Iteration 1 ended, frames my course in terms that map onto the four-theme architecture. The article's title is "Could AI be the next college teaching assistant? Some Colorado professors believe so." The framing is squarely within the Education theme. The article also reports on my course's industry-application orientation and on its accessibility implications, both of which map onto the other themes.

External journalism is not coached. The article's framing is the journalist's own organization of the material I had been teaching, and the framing lands on the architecture I had built. I treat this as additional cross-context evidence that the framework speaks intelligibly to non-academic audiences.

The article carries named quotes of me and of Ashley Stafford; I cite the article at source-level (AP-2024-05-16) in this dissertation.

A.5.4 Online workshop attendees engage with the framework in their feedback

The Iteration 3 Luma feedback corpus (LF-3, with sub-IDs LF-3.R01 through LF-3.R29) carries twenty-nine evaluative responses, nine of which include text feedback. The text feedback

engages the framework substantively.

One four-star reviewer wrote (text preserved verbatim from the source):

> "I think it was a good discussion regarding how to use the different AI image generation tools. A course work based on some of the neural networks behind them could be a great one."

The response is a constructive critique that surfaces a depth-versus-breadth tension within the Education theme. The reviewer wanted more on the neural-network mechanics under the tools (a deeper Education-theme dive into model architecture) and offered that the workshop succeeded as a tool-walkthrough (the breadth orientation). The tension is the kind of finding that analytic autoethnography surfaces from cross-context feedback. I return to it in §C.4.6.

A.5.5 My own public-facing writing across channels re-articulates the framework

The framework re-appears in my own public-facing writing across the cross-iteration channels.

The Keep Up Newsletter (KN-EP1 on image generation, KN-EP2 on research tools, KN-EP3 on sound tools) uses topic-clusters that align with the framework's domain coverage. Episode 1 sits within the Industry and Education themes (image generation as a craft and as a tool). Episode 2 sits within the Education and Industry themes (research tools as productive aids). Episode 3 sits within the Industry and Ethics themes (sound tools and the question of musician displacement).

The Keep Up Podcast (KP-EP2, KP-EP3) re-articulates the same theme clusters in audio form, building on the newsletter's groundwork.

The CU RMACC webinar (WB-2026-03-03), delivered ten weeks before my defense via the federal NAIRR Pilot platform, is the most recent public-facing synthesis of my work. It is also organized around the four-theme architecture, presenting the framework to a federal research audience as the coherent organizing principle of my pioneering practice.

A.5.6 What the cross-context evidence supports

The four-theme architecture is not a researcher's analytic imposition. It is the framework I built into my teaching from the K-12 origin (§A.2), and it is the framework that independently surfaces in five other sources: children at UW KidsTeam, undergraduates in my own courses applying it in final-project work and teach-outs, journalists at Aspen Public Radio, online attendees at the Iteration 3 workshop, and federal research audiences via the RMACC webinar.

I claim modest construct validity for the framework on this basis. The framework is not the only valid lens on generative-AI pedagogy. It is, however, a framework that has held its shape across age groups, institutional contexts, delivery modes, and audience types. Chapter D (§D.4) develops the framework's presence across eight documented contexts as a substantive finding in its own right.

Appendix B

Autoethnographic Supplement

B.1 Literature survey

The dissertation operates under two complementary methodological framings. The **primary** methodology, established in the proposal and presented in the main document, is design-based research (McKenney and Reeves 2018) grounded in constructivist learning theory and self-determination theory (Deci and Ryan 2000). It supports the iterative-refinement structure of the four course iterations and surfaces the three curriculum-design principles (modularity, learner choice, continuous feedback) that the main document develops.

This appendix introduces a **supplementary** methodological posture: analytic autoethnography (Anderson 2006). Analytic autoethnography is not a replacement for design-based research; it is an additional analytic lens that treats my position as the complete-member instructor of the four iterations as a resource for surfacing theoretical findings the DBR analysis developed in the main document does not surface. The two methodologies operate at different levels: DBR structures the empirical work of designing and revising the curriculum; analytic autoethnography surfaces what an instructor at the early-entrant moment of generative-AI pedagogy can analytically claim about her own practice.

This chapter surveys the autoethnography family, explains why analytic autoethnography supplements (rather than replaces) the primary methodology, and lays out how Anderson's five criteria will be applied to evidence the supplementary posture.

B.1.1 The autoethnography family

Autoethnography emerged in the 1990s as a research genre in which the researcher is also the subject of inquiry. Reed-Danahay (1997) edited the foundational volume **Auto/Ethnography: Rewriting the Self and the Social** and argued that the genre is best understood not as a single method but as a family of practices unified by the researcher's self-implication. Chang (2008) consolidated the methodological literature in **Autoethnography as Method**, naming procedures for data collection (personal memory, self-observation, self-reflection, and the use of external data such as artifacts and interviews) and for analysis (categorical aggregation, theming, narrative reconstruction).

Within this family, two principal variants emerged. **Evocative autoethnography**, articulated most fully by Ellis, Adams, and Bochner (2011) in their **Forum: Qualitative Social Research** overview, treats the autoethnographer as both author and subject of a narrative whose value lies in its affective resonance, its vulnerability, and the reader's identification with the lived experience portrayed. **Analytic autoethnography**, articulated by Anderson (2006) in **Journal of Contemporary Ethnography**, treats the autoethnographer as a complete member of a social setting whose insider position produces theoretical insight beyond what an external ethnographer could access. Anderson's variant retains the analytic posture of conventional ethnography while licensing the researcher's own experience as legitimate data.

Other variants in the family include collaborative autoethnography (Chang, Ngunjiri, and Hernandez 2013), in which multiple researchers analyze shared experience; narrative ethnography (Tedlock 1991), which foregrounds the researcher's narrative voice within an ethnographic project; and critical autoethnography, which pairs personal narrative with political critique. These variants matter as context but are not the supplementary framework this appendix adopts.

B.1.2 Two analytic layers, working together

The dissertation's two methodologies operate at different analytic layers and produce different kinds of contribution. Design-based research, as the main document develops it, surfaces three curriculum-design principles (modularity, learner choice, continuous feedback) by tracking what curricular changes worked across the four iterations. Analytic autoethnography, as this appendix develops it, **elaborates** each of those three principles by treating my reflexive position on the same iterations as an analytic resource, surfacing the patterns within each principle that the DBR analysis does not surface. Both layers read the same corpus; what changes between them is the analytic question asked.

Anderson (2006) reserves analytic autoethnography for cases in which the researcher is a complete member of the social setting under study, has produced reflective and curricular data during the practice itself, and is positioned to generate theoretical insight from the insider's position. Each of these preconditions describes my situation. The complete-member-researcher condition is satisfied: I was the instructor of record for Iterations 1 and 2 (CV-1, CV-2) and the workshop lead for Iterations 3 and 4 (DK-3, DK-4, TR-4.D1 through TR-4.D5). The reflective-data condition is satisfied: I produced fifteen weeks of contemporaneous reflective journaling in Iteration 1 (WU-1.W01 through WU-1.W15) and substantial public-facing reflective writing across the cross-iteration channels (RE, KN-EP series, KP-EP series, WB-2026-03-03). The theoretical-analysis condition is satisfied by the three nameable findings that emerge in Chapter D (hallucination-as-pedagogy, compression-as-curriculum-maturation, multi-channel teaching practice), which the DBR analysis developed in the main document does not surface.

Together, the two layers offer a more complete account of what the four iterations produced than either layer alone. The main document's principles are practitioner-facing outputs (what other curriculum designers can adopt); the appendix's findings are scholarship-facing outputs (what the field can claim about pioneer instructor practice). A reader of both encounters the principles in the main document and the findings here.

Of the autoethnography family, analytic autoethnography is the variant that fits this configuration. Evocative autoethnography (Ellis, Adams, and Bochner 2011) would call for sustained narrative immersion in the affective dimensions of teaching, which is a different kind of supplement and would sit awkwardly alongside the main document's DBR analysis. Analytic autoethnography's analytic posture is directly compatible with the DBR primary framing; both seek theoretical traction on what the iterations made possible to understand.

B.1.3 How I apply Anderson's five criteria

Anderson sets out five criteria for analytic autoethnography. Each maps onto a section of this methodology chapter and a body of evidence in my artifact corpus.

Criterion	Section	Evidence base
1. Complete member researcher status	§B.3.1	CV-1, CV-2, DK-3, DK-4, TR-4.D1 through TR-4.D5, HC series, KN-EP series, WB-2026-03-03
2. Analytic reflexivity	§B.3.2	WU-1.W01 through WU-1.W15, RE, KN-EP series
3. Narrative visibility of the researcher's self	§B.3.3	First-person voice across RE, KN-EP series, KP-EP series, WB-2026-03-03
4. Dialogue with informants beyond self	§B.3.4	23 named student outputs in DK-1.W01, LF-3.R01 through LF-3.R29, AP-2024-05-16, SP-2.* student teach-outs, KT-THEMES

Criterion	Section	Evidence base
5. Commitment to theoretical analysis	§B.3.5	Three nameable theoretical findings developed in Chapter D (hallucination-as-pedagogy, compression-as-curriculum-maturation, multi-channel teaching practice), which supplement the three curriculum-design principles the main document develops

Section B.3 applies each criterion in turn, evidences it against the corpus, and acknowledges where the evidence is thinner. The thinness acknowledgment is itself a move analytic autoethnography permits: criterion 2 (analytic reflexivity) is fully evidenced for Iteration 1 (the fifteen-week Weekly Updates Prelim Document is structured contemporaneous reflection) and partially evidenced for Iterations 2 through 4 through public-facing channels. I name this asymmetry rather than paper over it. Section B.3.2 characterizes the asymmetry precisely.

B.1.4 The relationship to the proposal's design-based research framing

The proposal my committee approved in October 2025 specified design-based research with thematic coding of interview data. The DBR portion of that framing has been retained in the main document; the iterative-refinement structure across the four iterations is the operational form DBR specified. What the proposal additionally specified (thematic coding of interview transcripts under IRB-governed conditions) was not enacted because the iterations did not produce the structured interview data thematic coding requires. The main document accordingly works from the artifacts the iterations did produce (curricular materials, surveys, workshop feedback) rather than from coded interviews.

This appendix's supplementary autoethnographic posture occupies the analytic space the

unrealized thematic-coding work would have occupied. Rather than coding student interviews to surface themes, I take the instructor's reflexive analytic posture on my own practice and the rich artifact corpus that practice generated. The two paths to theoretical insight are different in their data and procedures, but they are alike in their goal: drawing analytic claims from the iterations beyond the descriptive level the curriculum-design principles operate at.

Section B.2 develops how this supplementation works in practice.

B.2 How analytic autoethnography supplements design-based research

This section develops the supplementation relationship between the main document's primary methodology (design-based research with constructivist and self-determination-theory grounding) and this appendix's supplementary autoethnographic posture. The two methodologies operate at different levels and produce different kinds of contribution. Reading them together gives a fuller account of what my four iterations made possible to understand than either methodology alone.

B.2.1 What design-based research contributes

Design-based research, as McKenney and Reeves (2018) articulate it, is a methodology for designing educational interventions through iterative cycles in which the researcher specifies design conjectures, enacts them in real instructional settings, observes outcomes, and refines the conjectures across iterations. The main document of the dissertation operates under this methodology. The four course iterations are the iterative-design cycles; the curriculum-design principles (modularity, learner choice, continuous feedback) are what the DBR analysis surfaces; the comparison across the four iterations supports the principles' generalizability claims.

DBR is appropriate to this work. It licenses the iterative-refinement structure the iterations actually took; it surfaces principles that are practically useful to other curriculum designers; and it integrates the constructivist learning theory and self-determination theory that inform the curriculum's pedagogical design. The main document develops the DBR analysis at length.

B.2.2 What analytic autoethnography adds

What the DBR analysis developed in the main document does not surface is theoretical insight drawn from the instructor's reflexive analytic position on her own practice. Curriculum-design principles are practical outputs; they tell other curriculum designers what to do. They do not tell other researchers what the iterations reveal about how generative-AI pedagogy emerged as a field, how a pioneer instructor's practice operated under tool turnover, or what conceptual phenomena (like hallucination-as-pedagogy) the iterations exposed.

Analytic autoethnography (Anderson 2006) is the methodological framework that licenses these additional analytic claims. By treating my insider position as an analytic resource rather than as a confound to be controlled, autoethnography surfaces:

- **Theoretical findings the DBR analysis does not produce.** Chapter D develops three: hallucination-as-pedagogy, compression-as-curriculum-maturation, multi-channel teaching practice. These are not curriculum-design principles for adopters; they are analytic claims about the pioneer-instructor practice at the generative-AI moment.
- **Reflexive accounting for my position.** Chapter B §B.7 performs reflexivity on my K-12-to-CU trajectory, my investment in the work, and the methodological-supplementation choice itself.
- **Use of a richer artifact corpus.** The Weekly Updates Prelim Document, the Keep Up Newsletter, the Keep Up Podcast, the CU RMACC webinar, the UW KidsTeam co-design data, the Aspen Public Radio coverage, and the AI-IRT Seed Grant proposal are autoethnographic data sources that the DBR analysis does not naturally read for theoretical content.

The supplementation is additive. The DBR analysis of the main document is unaffected by what this appendix does; what this appendix does is sit alongside it and extend the analytic reach.

B.2.3 What the iterations produced for both methodologies

The four iterations between Spring 2024 and September 2025 produced data of two kinds simultaneously. Curricular and operational data fed the DBR analysis: weekly module schedules, slide decks, assignments, the Luma feedback survey for Iteration 3, attendance and engagement metrics, syllabi. Reflective and contextual data fed (and feeds) the autoethnographic analysis: the Weekly Updates Prelim Document for Iteration 1, the cross-iteration newsletter and podcast, the CU RMACC webinar, the K-12 outreach corpus, external media coverage.

The two streams of data are not separate corpora; they overlap. The Canvas LMS exports (CV-1, CV-2) feed both analyses: they document curriculum design (DBR-relevant) and they document the iteration in the form a complete-member researcher can read autoethnographically. The slide decks across iterations (DK-1.FG, DK-1.W01, DK-2 series, DK-3, DK-4) are similar. What changes between the two methodologies is the analytic question asked of the same artifacts: "what curriculum-design principle does this iteration surface?" (DBR) versus "what theoretical claim does my reflexive position on this iteration license?" (analytic autoethnography).

B.2.4 Why analytic autoethnography is the right supplementary methodology

I considered three candidate supplementary methodologies before choosing analytic autoethnography. Each of the others would have offered something but did not fit the configuration as well.

Thematic coding of student interviews was the supplementary path the original proposal specified. It was not enacted because the iterations did not produce IRB-governed structured interview data. Other learner-facing data (the Luma feedback corpus, the student teach-out presentations, the UW KidsTeam co-design data) was generated, and the main document's DBR analysis reads it; but the data does not match what thematic coding presumes.

Evocative autoethnography (Ellis, Adams, and Bochner 2011) would have produced an affectively resonant first-person narrative of the teaching practice. The genre is established and well-developed. It would, however, sit awkwardly alongside the main document's DBR analysis;

evocative narrative and DBR analysis operate from different epistemological commitments. The combined dissertation would read as two unrelated works rather than as a coherent supplementation.

Practitioner inquiry (Cochran-Smith and Lytle 2009) is closer in spirit to what this appendix does, and the appendix draws on it as background. As a primary supplementary methodology, however, practitioner inquiry typically operates at the same level as DBR (surfacing practical knowledge for teachers) and would therefore not add an additional analytic layer.

Analytic autoethnography offers what the supplementation requires: a methodology with its own rigor structure (Anderson's five criteria), its own established place in engineering education and HCI literatures, its own genre of contribution (theoretical findings from insider analytic position), and full compatibility with the DBR primary framing of the main document. Chapter B §B.3 below applies Anderson's criteria to the corpus.

B.2.5 What this supplementation does not claim

The supplementation does not claim that analytic autoethnography is necessary for the dissertation. The main document stands on its DBR analysis; the curriculum-design principles do not require autoethnographic support to be valid. What the supplementation claims is that the autoethnographic posture surfaces an additional layer of theoretical findings the DBR analysis developed in the main document does not surface, and that those findings are themselves valuable scholarly contributions.

The supplementation also does not claim methodological neutrality. Section B.7 names my investment in the work. The supplementation's purpose is to enrich the analytic record, not to provide a second independent test of the main document's claims.

With this supplementation framing established, §B.3 below applies Anderson's five criteria to my artifact corpus in detail.

B.3 Anderson's five criteria, each evidenced

Anderson (2006) sets out five criteria that distinguish analytic autoethnography from neighboring genres and govern its rigor. In this section I take each criterion in turn, state what Anderson requires, and evidence the criterion against my artifact corpus.

B.3.1 Complete member researcher status

Anderson requires that the autoethnographer be a complete member of the social setting under study, not a peripheral participant or a visiting researcher. The framework draws on Adler and Adler's (1987) typology of membership roles and reserves analytic autoethnography for the case in which the researcher is fully embedded in the setting and has a stake in it.

I was the complete member of every teaching setting documented in my dissertation. For Iterations 1 and 2 I was the instructor of record (CV-1, CV-2, SY-2). For Iterations 3 and 4 I was the workshop lead and primary deliverer (DK-3, DK-4, TR-4.D1 through TR-4.D5). For the HCI summer 2024 guest-lecture series I was the invited instructor across ten sessions (HC-INTRO, HC-ACCESS, HC-AUDIO, HC-EDU, HC-INDUSTRY, HC-MUSIC, HC-VIDEO-ETH, HC-VIDEO, HC-FUTURE-WHEEL, HC-SORA). For my LinkedIn newsletter and podcast I was the author and host (KN-EP1 through KN-EP3, KP-EP2, KP-EP3). For the CU RMACC webinar delivered through the federal NAIRR Pilot program I was the named speaker (WB-2026-03-03). I delivered the AI Art Mural at the CU STEAM Festival (ST-MURAL, ST-PHOTO) as the workshop lead, and I led the K-12 art contest documented in my Research Impact Essay (RE).

Complete member researcher status is not merely a credentialing claim. It conditions what I could learn and how I could learn it. Because I was the instructor, I could revise the curriculum in real time on the basis of what I was observing (the tool turnover and module reshuffles documented in §C.6.1 and §C.6.2 are the visible trace of this). Because I had a stake in the work, I had reason to write fifteen weeks of structured reflective notes in Iteration 1 (WU-1.W01 through WU-1.W15) and to produce public-facing reflective writing across the cross-iteration channels. The data I have

is the data only a complete member could have generated.

B.3.2 Analytic reflexivity

Anderson requires that the autoethnographer practice reflexivity, meaning sustained self-aware reflection on the researcher's position, choices, and learning. Anderson distinguishes this from confessional or therapeutic narrative; reflexivity in analytic autoethnography is analytic, not catharsis.

The Weekly Updates Prelim Document is my central piece of evidence for this criterion. Across fifteen weeks of Iteration 1, I logged what I had taught, who had guest-spoken, what resources I had used, what I had learned, and what students had produced. The "Learned" entries in particular function as analytic reflexivity: each week I named what the teaching had taught me. In Week 1 (WU-1.W01-Q1) I wrote "I learned that some of the multiple choice quizzes generated by ChatGPT were not correct and had hallucinations." That observation, made in the first week of the first iteration, becomes the seed for the hallucination-as-pedagogy finding I develop in §D.2.

Beyond the Weekly Updates, my Research Impact Essay (RE) is a piece of first-person reflective writing produced near the end of Iteration 1. Its narrative arc — from a school-district professional development through a high-school art contest to the pilot CU course (RE-Q3) — is the explicit reflexive account of how I came to be doing this work. Three episodes of my Keep Up Newsletter (KN-EP1, KN-EP2, KN-EP3) and two episodes of my podcast (KP-EP2, KP-EP3) extend the reflexive voice into public-facing media in 2025.

I acknowledge an important asymmetry here, and characterize it precisely. The structured weekly reflection I kept in Iteration 1 (WU-1.W01 through WU-1.W15) is contemporaneous, structured, and analytic in the sense Anderson requires: each weekly "Learned" entry names what the teaching had taught me, captured at the time the teaching happened. The cross-iteration channels (KN-EP series, KP-EP series, WB-2026-03-03) carry reflective content but are not equivalent. They are public-facing writing, shaped by audience and platform; they were produced months or years after the iterations they describe; and they serve a presentational function in addition to a

reflective one. Calling them "analytic reflexivity" in Anderson's strict sense overstates what they are.

What I claim, more precisely, is this: Iteration 1 satisfies Anderson's analytic-reflexivity criterion fully. Iterations 2, 3, and 4 are supplemented by public-facing reflective writing that I acknowledge is retrospective and presentational rather than contemporaneous and analytic. Analytic autoethnography permits this supplementation when the retrospective-public frame is named (as I name it here) rather than concealed, but the supplementation is partial. The reflective base for the later iterations is genuinely thinner than for Iteration 1, and the findings developed in Chapter D rely accordingly on cross-source triangulation (§B.6.2) rather than on per-iteration reflective depth.

B.3.3 Narrative visibility of the researcher's self

Anderson requires that the researcher's self be visible in the text, not hidden behind a falsely-detached scholarly voice. The autoethnographer must appear as a character in the narrative, with named actions, named choices, and named consequences.

I am visible throughout my artifacts. My Research Impact Essay is written in the first person from start to finish (RE). My Keep Up Newsletter episodes use the running-and-training metaphor to narrate my own ongoing engagement with new generative-AI tools (KN-EP1 through KN-EP3). My CU RMACC webinar is delivered as a personal account of what I have built (WB-2026-03-03). The ForeverGold deck for Iteration 1 (DK-1.FG) names me by affiliation (Imagine AI Lab, advised by Prof. Tom Yeh) on the title slide. The dissertation itself is written in the first person, and this revised methodology chapter is a more explicit performance of the narrative-visibility criterion than the original draft.

Narrative visibility is also visible in what I name as my position. I am a former English-language-arts and high-school art teacher (RE) who entered a PhD program at the College of Engineering, lab-affiliated with the ATLAS Institute, and crossed into the College of Engineering and Applied Science for Iteration 3. My position as someone whose disciplinary home is not computer science but who teaches generative-AI software practice to engineering undergraduates

is a positionality finding in its own right, developed in §B.3 of Chapter A.

B.3.4 Dialogue with informants beyond self

Anderson requires that the autoethnographer engage with informants beyond the self. The framework is not solipsistic; the researcher's insider voice must converse with the voices and observations of others in the setting.

My artifact corpus carries dialogue beyond self at several scales.

Students in my own courses. The Week 1 deck for Iteration 1 (DK-1.W01) contains twenty-three named student prompt-engineering outputs from the opening session. Iteration 2 generated six student teach-out presentations preserved in the archive (SP-2.DEEPFAKE, SP-2.SINTRA, SP-2.HALEY, SP-2.DAKOTA, SP-2.ROBOTICS). The Iteration 1 Final Project Requirements (FP-1) document the Media Studies company-creation assignment that surfaced students' own four-theme syntheses.

Online learners. The Luma feedback corpus for Iteration 3 (LF-3) carries twenty-nine evaluative responses with text feedback in nine of them. The participant roster for the same iteration (LR-3) documents the audience composition: 411 registered, 129 attended live, sixty-five percent students, seventy percent expressing interest in a Master's program.

Children in K-12 settings. The UW KidsTeam research collaboration (KT-DECK, KT-IDEAS, KT-COMIC, KT-NOTES, KT-THEMES, KT-YAB) surfaced child and teen voices on generative AI in schools across a three-day session and multiple YAB sessions in July 2024. KT-THEMES catalogues the challenges and opportunities the children themselves named, including the "third arm" hallucination observation (KT-THEMES-C5) that independently mirrors my own first-week observation in Iteration 1.

External commentators. The Aspen Public Radio article (AP-2024-05-16) published in May 2024 quotes both me and my Iteration 1 student Ashley Stafford on what the course offered.

Co-instructors and guest speakers. Across both undergraduate iterations the guest-speaker roster operated as a curated dialogue with industry, ethics, and adjacent-research voices:

Anthony Pinter, Diane Sieber, Daniel Ritchie, Nikolaus Klassen, Suibi Weng, and Nolan Brady in Iteration 1; Tom Yeh, Nolan Brady, Bobby Hodgkinson, and Justin Shacklette in Iteration 2.

Dialogue with informants beyond self is not a list of names. It is the evidentiary base for the cross-iteration findings I develop in Chapter D. Children at UW KidsTeam observing the same hallucination phenomenon I observed in Iteration 1 (KT-THEMES-C5 and WU-1.W01-Q1) is what makes hallucination-as-pedagogy a triangulated finding rather than an instructor's stray observation.

B.3.5 Commitment to theoretical analysis

Anderson's fifth criterion is the one most distinctive of analytic autoethnography. The autoethnographer must use the insider position not to produce evocative description but to generate, refine, or extend theoretical understanding of broader social phenomena. Anderson is explicit that this criterion is the hardest to meet and the most often missing in autoethnographic work that drifts toward the evocative.

I commit to three theoretical claims, each developed at length in Chapter D and each cross-validated against multiple sources in the corpus.

Hallucination as pedagogy. The phenomenon of generative-AI hallucination is most commonly framed in the technical and policy literature as a system limitation to be reduced or guarded against (Bender et al. 2021 is the canonical formulation). My data suggests a different framing for the pedagogical setting: hallucination is a productive teachable moment that surfaces both the limits of the technology and the importance of human verification. I draw this claim from four independent sources: my own Iteration 1 Week 1 reflection (WU-1.W01-Q1), the UW KidsTeam children's independent observation of "third arm" images (KT-THEMES-C5), my public-facing Keep Up Newsletter framing of hallucination as expected behavior (KN-EP1-Q1), and the live workshop delivery in Iteration 4 (TR-4.D1 carries hallucination as a teaching topic).

Compression as curriculum maturation. Across the four iterations my curriculum compressed from fifteen weeks to five days while retaining the same four-theme architecture. The

compression ratio is approximately six to one. My theoretical claim is that this represents curriculum maturation by distillation rather than by accretion, and that pioneering instructor practice in fast-moving technology fields produces this pattern naturally. The claim draws on the iteration comparisons developed in §C.6 of Chapter C.

Multi-channel teaching practice. My pioneering practice did not unfold in a single classroom. It unfolded across eight documented channels: two undergraduate semester courses, two online workshops, an HCI graduate guest-lecture series, K-12 outreach activities, a public-facing LinkedIn newsletter and AI by Hand podcast, and a federal research webinar. My theoretical claim is that practitioner-pioneer practice at the technological frontier is most accurately characterized as networked multi-channel engagement rather than as discrete classroom delivery, and that this characterization has implications for how engineering education and HCI scholarship treats pioneering teaching work. The channels are catalogued in §D.4.

Each of these claims is a theoretical contribution in the sense Anderson specifies: a generalizable analytic point that the insider position produced and that the artifact corpus supports. Chapter D develops each claim against its evidence in detail.

B.3.6 Summary

Anderson's five criteria are met by what I did and what data I have. Complete member researcher status: I was the instructor at every site. Analytic reflexivity: I kept structured reflective notes in Iteration 1 and produced extensive public-facing reflective writing across the iterations. Narrative visibility: my first-person voice is present across the corpus. Dialogue beyond self: I have student work, learner feedback, child-research-participant accounts, and external media coverage. Commitment to theoretical analysis: I make three nameable theoretical claims, each cross-validated across multiple sources.

The methodology fits the work, and the work fits the methodology.

B.4 Data sources

My data set is organized in five categories. In this section I name each category, indicate what kind of material it covers, and point the reader to Appendix E for the per-artifact catalog. Volume figures throughout are summary numbers; specific IDs, dates, and descriptions live in the catalog.

B.4.1 Curricular artifacts

Curricular artifacts are the materials I produced or organized as the instructor: Canvas LMS exports for the two semester iterations (CV-1, CV-2), the Iteration 2 syllabus (SY-2), the ForeverGold course deck for Iteration 1 (DK-1.FG), the Week 1 opening lectures and other iteration-specific lecture decks (DK-1.W01, DK-2.JAN13 through DK-2.MAR10), the workshop decks for Iterations 3 and 4 (DK-3, DK-4), the Iteration 1 Final Project assignment (FP-1), the ten HCI summer 2024 guest-lecture decks (HC corpus), the DLS Prompt Engineering deck (DK-DLS), and the K-12-oriented Storytelling Cartoonimator worksheet (STC). The two Canvas exports together hold 465 files of week-by-week curricular structure including assignments, attendance, readings, rubrics, and ten video recordings.

B.4.2 Reflective artifacts

Reflective artifacts are the materials in which I narrated my own learning and choices as the instructor. The fifteen-week Weekly Updates Prelim Document (WU-1.W01 through WU-1.W15) is the densest single reflective source and the only contemporaneous structured journal in the corpus. The Research Impact Essay (RE) provides first-person reflective writing from near the end of Iteration 1. Three episodes of the Keep Up Newsletter on LinkedIn (KN-EP1 through KN-EP3) and two episodes of the Keep Up Podcast (KP-EP2, KP-EP3) carry my public-facing reflective voice across 2025. The CU RMACC webinar (WB-2026-03-03) delivered through the federal NAIRR Pilot platform is the most recent public-facing synthesis of my work, ten weeks

before defense. Combined reflective text is approximately twenty thousand words of first-person material across the cross-iteration channels, plus the structured fifteen weekly entries for Iteration 1.

B.4.3 Delivery artifacts

Delivery artifacts are the verbatim records of what I said as I taught. The principal delivery corpus is the set of five day-by-day transcripts from Iteration 4 (TR-4.D1 through TR-4.D5), totaling approximately 55,000 words of recorded teaching. The Canvas exports for Iterations 1 and 2 contain ten additional guest-lecture video recordings (four in CV-1, six in CV-2); those recordings are part of the corpus by reference but are not used as text-level evidence in this analysis.

B.4.4 Learner-facing artifacts

Learner-facing artifacts are the materials produced by students, online learners, child research participants, and external commentators in response to my teaching. They include the 23 named student outputs displayed in the Iteration 1 Week 1 deck (within DK-1.W01), the six Iteration 2 student teach-out presentations (SP-2 series), the 29 Iteration 3 Luma feedback responses with nine carrying text comments (LF-3 and its R-numbered sub-IDs), the Iteration 3 Luma participant roster of 411 registered and 129 live attendees (LR-3), the UW KidsTeam co-design research with children and teens (KT-DECK, KT-IDEAS, KT-COMIC, KT-NOTES, KT-THEMES, KT-YAB), the STEAM Festival mural and accompanying photograph (ST-MURAL, ST-PHOTO), and the Aspen Public Radio article that quoted me and my Iteration 1 student Ashley Stafford one week after Iteration 1 ended (AP-2024-05-16).

B.4.5 External artifacts and program documents

External artifacts situate my work within institutional, funding, and program contexts. They include the AI-IRT Seed Grant proposal whose principal investigators are two of my committee members (AI-PROPOSAL), and the ENED Preliminary Exam Part 2 document that records my

commitment to offering a continuous generative-AI course at CU Boulder (PR-PART2).

B.4.6 Categorization summary

The data set is contemporaneous, instructor-present, and substantial. Iteration-specific text data is approximately 105,000 words across the four iterations; cross-iteration material adds approximately 35,000 words of first-person and public-facing writing, plus the ten unprocessed video recordings and the K-12 and grant documents. Total contemporaneous and instructor-present text data exceeds 140,000 words. This is well above the empirical saturation threshold for analytic autoethnography.

I acknowledge two asymmetries. First, my contemporaneous reflective journaling is concentrated in Iteration 1. Iterations 2, 3, and 4 draw on the cross-iteration reflective channels (KN-EP series, KP-EP series, WB-2026-03-03) as the supplementary reflective base, within an acknowledged retrospective-public frame. Second, ten guest-lecture video recordings in CV-1 and CV-2 are part of the corpus by reference but are not used as text-level evidence in the analysis. I name both asymmetries explicitly rather than claiming a more uniform evidentiary surface than the data supports. Section B.5 describes the analytic process I apply to this corpus.

B.5 Analytic process

The analytic process I apply to my corpus has four components: narrative reconstruction per iteration, cross-iteration pattern matching, anchor-concept development, and member checking. Each is described below.

B.5.1 Narrative reconstruction per iteration

For each of my four iterations I produce a narrative reconstruction grounded in the artifacts specific to that iteration. The reconstruction follows the structure analytic autoethnography requires: it is chronological where chronology matters, it is sourced to specific artifacts, and it foregrounds my position as the instructor making curricular and pedagogical choices. The recon-

struction also draws on Schön's (1983) framing of the reflective practitioner: the practice's choices are reconstructed from artifacts that contain "reflection-in-action" traces (the Weekly Updates Prelim Document's "Learned" entries), not retrospectively imposed on the iterations from the outside.

For Iteration 1, the reconstruction draws on the Canvas LMS export (CV-1) for week-by-week curricular structure, on the Weekly Updates Prelim Document (WU-1.W01 through WU-1.W15) for my contemporaneous reflective voice, on the ForeverGold deck (DK-1.FG) for the course-framing architecture, on the Week 1 deck (DK-1.W01) for the opening session's named student outputs, and on the Final Project Requirements (FP-1) for the capstone assignment. The reconstruction is supplemented by the Research Impact Essay (RE), which I wrote near the end of Iteration 1.

For Iteration 2, the reconstruction draws on the Canvas export (CV-2), the syllabus (SY-2), the five iteration-specific decks (DK-2.JAN13 through DK-2.MAR10), the six student teach-out presentations (SP-2.**), and the administrative spreadsheets (AS-2.**) that document the iteration's guest-lecturer list, teach-out schedule, and final-presentation order.

For Iteration 3, the reconstruction draws on the workshop deck (DK-3), the Luma feedback corpus (LF-3 with sub-IDs LF-3.R01 through LF-3.R29), and the participant roster (LR-3).

For Iteration 4, the reconstruction draws on the workshop deck (DK-4) and the five day-by-day YouTube transcripts (TR-4.D1 through TR-4.D5), which together provide approximately fifty-five thousand words of verbatim teaching delivery.

Each narrative reconstruction is presented in Chapter C (§C.2 through §C.5).

B.5.2 Cross-iteration pattern matching

After reconstructing each iteration individually I conduct a cross-iteration pattern-matching pass. The goal is to surface the regularities and the changes that running the same course four times across different sites and durations exposes.

The pattern-matching pass produces specific findings of the following kinds:

Tool turnover. Comparison of CV-1 against CV-2 produces a documented list of tools added (DeepSeek added in Iteration 2 Week 3 within weeks of its January 2025 public release; AI

Agents as a full new module Week 9; NotebookLM moved from student-mentioned in Iteration 1 to dedicated lecture in Iteration 2; Be My AI, Sora, HeyGen, ElevenLabs, Custom GPTs with Wolfram API), tools repositioned (Hugging Face moved from Iteration 1 Week 9 to Iteration 2 Week 7), and tools dropped (Reality Editor present in Iteration 1 Week 10, absent in Iteration 2). The full list is presented in §C.6.1.

Module placement evolution. Same comparison surfaces specific structural shifts: the video module moved from Iteration 1 Week 6 to Iteration 2 Week 14, an eight-week shift; the industry theme moved from Iteration 1 Week 5 to Iteration 2 Week 8. The full list is presented in §C.6.2.

Guest-speaker turnover. Cross-comparison of the guest-speaker rosters surfaces a clean pattern: external guests turn over completely between Iterations 1 and 2 (Pinter, Sieber, Zago, Klassen, Weng all in Iteration 1 but not Iteration 2), while CU-internal speakers are stable but rotate topics (Nolan Brady returns from Iteration 1 to Iteration 2 with a different lecture). This pattern is documented in §C.6.3.

Compression evidence. Comparison of the four iterations on time-on-task produces the headline compression finding: fifteen weeks for Iterations 1 and 2, five days at one hour per day for Iterations 3 and 4. The same four-theme architecture survives both formats. The compression ratio is approximately six to one.

Stable elements. Cross-comparison also surfaces what does not change. The Week 1 prompt-engineering opener is identical in structure across all four iterations. The four-theme architecture is present in every iteration's framing. The Midjourney Self-Portrait assignment recurs.

B.5.3 Anchor-concept development

From the cross-iteration patterns and from my reflective material I develop three anchor concepts, each of which becomes a theoretical claim in Chapter D. The development process is iterative and reflexive: I name a candidate concept, I test it against the artifact corpus, I refine the concept's wording, and I retain it as an anchor only if it withstands sourcing against multiple

independent artifacts.

Hallucination-as-pedagogy anchors my first finding. The concept names a reframing of generative-AI hallucination from a system limitation to a teachable pedagogical moment. I tested the concept against four independent sources: my own Iteration 1 Week 1 reflection (WU-1.W01-Q1, "I learned that some of the multiple choice quizzes generated by ChatGPT were not correct and had hallucinations"), the UW KidsTeam children's independent surfacing of "images produced with a third arm" as a concern (KT-THEMES-C5), my public-facing Keep Up Newsletter framing of hallucination as expected behavior (KN-EP1-Q1, "Expect variable results and occasional hallucinations"), and the live workshop delivery in Iteration 4 (TR-4.D1 carries hallucination as a teaching topic). The concept survived testing and is retained.

Compression-as-curriculum-maturation anchors my second finding. The concept names a pattern in which the same content architecture survives a six-fold compression from a fifteen-week semester to a five-day workshop, and in which the compression represents maturation by distillation rather than content loss. I tested the concept against the four iteration artifact sets and against my reflective voice across the cross-iteration material. The concept survived testing and is retained.

Multi-channel teaching practice anchors my third finding. The concept names a configuration in which pioneer instructor practice unfolds simultaneously across multiple delivery channels rather than within a single classroom. I tested the concept against the documented channels: two undergraduate semester courses (CV-1, CV-2), two online workshops (DK-3 with LF-3, DK-4 with TR-4.D5), an HCI grad guest-lecture series (HC series), K-12 outreach (ST-MURAL, KT-DECK), a LinkedIn newsletter (KN-EP series), a podcast (KP-EP series), and a federal research webinar (WB-2026-03-03). All eight channels are documented; the concept survived testing and is retained.

Section D.2, 4.3, and 4.4 develop each anchor concept into a theoretical claim.

B.5.4 Member checking by chair and committee

Analytic autoethnography permits and encourages member checking with collaborators and advisors who have direct knowledge of the setting. My chair has been present in the setting in

two ways. He is co-PI on the AI-IRT Seed Grant that funded the research arc (AI-PROPOSAL), and he delivered the DeepSeek guest lecture in Iteration 2 (named in CV-2 and on the Iteration 2 guest-lecturer spreadsheet AS-2.GUESTS, and appearing live in the Iteration 4 Day 1 transcript TR-4.D1 as a guest from CU Boulder). His member-checking role is not extra-textual; it is partially documented in the corpus itself.

My second committee member, Diane Sieber, is also co-PI on the AI-IRT Seed Grant and delivered the writing-with-GenAI guest lecture in Iteration 1 (named in CV-1). Her member-checking role is similarly partially documented.

I treat the supervisory and committee feedback I have received during proposal review and during the dissertation submission process as part of the analytic process rather than as external editorial input. This is consistent with Anderson's framework, which positions the researcher's analytic position as informed by but not subordinate to the surrounding intellectual community.

B.5.5 Artifact provenance and dating

Every artifact cited in this appendix carries provenance: a date or date range and a position within the iteration framework or cross-iteration timeline. The full per-artifact catalog appears as Appendix E, organized by iteration and grouped by artifact category, with brief descriptive entries for every ID the chapters cite.

I use three dating conventions throughout. **Exact dates** when the artifact is precisely dated (the Aspen Public Radio article AP-2024-05-16 is dated May 16, 2024; the CU RMACC webinar WB-2026-03-03 was delivered March 3, 2026). **Approximate dates** marked with `c.` when an internal cue locates the artifact within a season but not a specific day (the Research Impact Essay is dated c. Spring 2024 from its "during this spring semester" language). **Undated** when no internal cue is available (the DLS Prompt Engineering deck DK-DLS carries no internal date marker; Appendix E flags it as undated).

Provenance integrity is the basis for the trustworthiness claims developed in §B.6. A reader can audit any in-text citation against Appendix E to confirm the artifact's date, iteration position,

and descriptive content.

B.6 Trustworthiness

Trustworthiness in analytic autoethnography is established through Anderson's analytic rigor criteria, through triangulation across artifact types and contexts, and through honest acknowledgment of the limits of the data and of my position as the researcher.

B.6.1 Anderson's analytic rigor as the trustworthiness frame

Section B.3 lays out Anderson's five criteria and evidences each against my corpus. The five criteria are themselves the trustworthiness frame for analytic autoethnography. A work is trustworthy as analytic autoethnography to the extent that it satisfies the criteria; a work that drifts toward evocative narrative without the criteria's analytic discipline is a different kind of work and is evaluated differently.

My corpus satisfies the criteria as evidenced in §B.3. Complete member researcher status is documented across every site of teaching practice (CV-1, CV-2, DK-3, DK-4, the HC series, the KN series, the KP series, WB-2026-03-03). Analytic reflexivity is most fully evidenced in the Iteration 1 Weekly Updates Prelim Document (WU-1.W01 through WU-1.W15) and is supplemented by my public-facing reflective writing across the cross-iteration channels. Narrative visibility of my self is present throughout the corpus, with first-person voice in the Research Impact Essay (RE), the newsletter (KN-EP series), the podcast (KP-EP series), and the webinar (WB-2026-03-03). Dialogue with informants beyond self is evidenced by twenty-three named student outputs in DK-1.W01, twenty-nine Luma responses (LF-3.R01 through LF-3.R29), the Aspen Public Radio coverage (AP-2024-05-16), the UW KidsTeam research data (KT-THEMES and the broader KT corpus), six student teach-out presentations (SP-2.*), and the documented guest-speaker rosters across both iterations. Commitment to theoretical analysis is satisfied by the three anchor concepts I develop in Chapter D (hallucination-as-pedagogy, compression-as-curriculum-maturation, multi-channel teaching practice), each tested against multiple independent sources.

B.6.2 Triangulation across artifact types and contexts

The strongest single trustworthiness move I make is triangulation: anchoring each substantive claim in multiple independent artifacts of different types from different points in the timeline.

My hallucination-as-pedagogy claim is anchored in four independent sources: my contemporaneous reflective journal (WU-1.W01-Q1, January 2024), an independent K-12 research data set (KT-THEMES-C5, July 2024), my public-facing reflective writing (KN-EP1-Q1, April 2025), and my recorded teaching delivery (TR-4.D1, September 2025). These four sources span roughly twenty months and three distinct types of artifact: instructor reflection, independent learner observation, public-facing writing, and recorded delivery. The hallucination phenomenon as pedagogical opportunity surfaces in all four.

My compression-as-curriculum-maturation claim is anchored in the four iteration artifact sets. Iteration 1 (CV-1, DK-1.FG, DK-1.W01, FP-1, WU-1 series) shows the fifteen-week structure with the four-theme architecture explicit. Iteration 2 (CV-2, SY-2, DK-2 series) shows the same architecture in a second fifteen-week iteration, with documented tool turnover and module reshuffles. Iteration 3 (DK-3, LF-3, LR-3) shows the architecture surviving compression to five days. Iteration 4 (DK-4, TR-4.D1 through TR-4.D5) shows the same five-day compression carrying through to a second cohort.

My multi-channel teaching practice claim is anchored across the eight documented channels described in §B.3.5 and developed in §D.4. Each channel is a separate corpus with its own artifact provenance.

Triangulation is the practical instrument by which analytic autoethnography moves from instructor's stray observation to defensible theoretical claim.

B.6.3 Honest acknowledgment of the limits of the data

I name two limits of my data that the methodology chapter must surface rather than paper over.

Asymmetry in contemporaneous reflective material. Iteration 1 is uniquely well-served by the fifteen-week Weekly Updates Prelim Document. Iterations 2, 3, and 4 do not have a structured contemporaneous reflective counterpart. The cross-iteration reflective channels (KN-EP series across the 2025 iterations, KP-EP series in May 2025, WB-2026-03-03 in March 2026) carry reflective writing across the later iterations within an acknowledged retrospective-public frame, and analytic autoethnography permits this configuration when the frame is named.

Aspen Public Radio quoted at source-level. The Aspen Public Radio article (AP-2024-05-16) is held as a source-level reference in the evidence-table. The article's named quotes of me and of my Iteration 1 student Ashley Stafford are part of the corpus by reference; the dissertation cites the article at source-level rather than at sentence-level.

These limits are not invalidating. They are common in analytic autoethnographic work that surveys a multi-year practitioner-pioneer practice across multiple delivery channels. Naming them is what trustworthy autoethnography requires.

B.6.4 Selection of the corpus itself

The corpus is a curated artifact. I chose what to include, in what level of detail, and with what framing. Selection is itself a methodological act that requires acknowledgment.

What I included: every artifact I produced as the instructor of the four iterations that I could locate and ingest (Canvas LMS exports, slide decks, syllabi, the Weekly Updates Prelim Document, the Final Project assignment); the publicly-available cross-iteration material I authored (Research Impact Essay, Keep Up Newsletter, Keep Up Podcast, CU RMACC webinar); learner-facing material that captures student or learner voice (named student outputs in DK-1.W01, the Iteration 2 student teach-out presentations, the Iteration 3 Luma feedback corpus); collaborative research material relevant to the K-12 strand (the UW KidsTeam corpus); and institutional context (the AI-IRT Seed Grant proposal, the ENED Preliminary Exam Part 2).

What I included only partially: the Aspen Public Radio article cited at source-level rather than sentence-level (§B.6.3), and feedback channels for Iterations 1 and 2 that would parallel the

Luma corpus for Iteration 3 (these were not collected through a structured instrument because the semester iterations did not run through Luma).

A different researcher with access to the same setting might have included different material. This is the honest position selection bias requires: not neutrality, which is unavailable, but transparency about what was included and what was not.

B.6.5 Identifiability and PII review scope

The autoethnographic appendix differs from the main document in the kinds of identifiable mention it contains. The main document analyzes the instructor-authored curriculum and treats course evaluation feedback as anonymized triggers for revision (main §2.6.5). The appendix, by contrast, is autoethnographic: it draws on the instructor's reflexive position, names participants in the four iterations where they appear in the artifact corpus, and quotes from public-facing materials (the Keep Up Newsletter, the Keep Up Podcast, the CU RMACC webinar, the Aspen Public Radio article). This wider corpus carries identifiable mentions that need to be accounted for separately.

Four categories of identifiable mention appear in this appendix, each handled differently:

- **Students named in instructor-produced course materials** (for example, the DK-1.W01 opening-deck student-output gallery; named Iteration 1 students quoted in external coverage such as AP-2024-05-16). These will be reviewed before final submission. Where written consent for educational use is documented, the name remains and the consent is noted; where consent is not documented, the mention will be anonymized.
- **Named guest speakers and co-instructors** across the four iterations and the HCI summer 2024 series. These are public professional identities — adults who agreed to deliver guest lectures and panels — and they are retained in the appendix with their professional context.
- **Workshop participants in the Luma feedback corpora** (LF-3.R01 through LF-3.R29 for Iteration 3; LR-4 and LF-4 for Iteration 4). Identifiable to the workshop organizer

through the Luma platform. Quotes used in the appendix will be anonymized in the published version.

- **External journalism and public-facing reflective writing** (AP-2024-05-16, KN-EP series, KP-EP series, WB-2026-03-03). Already published with the participants' consent through journalism, podcast, or webinar channels; retained as cited.

Handling of these categories has been and will remain consistent with the Family Educational Rights and Privacy Act (FERPA, 20 U.S.C. § 1232g; 34 CFR Part 99) and the University of Colorado Boulder's policies on student records. The appendix's autoethnographic framing follows the standards of analytic autoethnography (Anderson 2006), in which the instructor's reflexive position and named relationships in the practice are part of the analytic subject. As with the main document's human-subjects scope (main §2.6.5), the analytic claims of this appendix do not require IRB-reviewed measurement of student outcomes.

B.6.6 Honest acknowledgment of my position as the researcher

I am the subject of the work and the researcher of the work. This is the definitional condition of autoethnography, and it carries a definitional limitation: my interpretive choices are not neutral. I have a stake in my work succeeding, in my findings landing as contributions, and in the dissertation being awarded. Any reader of an analytic autoethnography must take this stake into account.

I respond to this limitation through three moves. First, I source every substantive claim to artifacts produced at the time of the events described, rather than to retrospective reconstructions, wherever the artifacts allow. The Weekly Updates Prelim Document is the strongest such anchor for Iteration 1. Second, I subject every anchor concept to multi-source testing as described in §B.5.3 and §B.6.2, rather than committing to a concept on the strength of a single observation. Third, I expose my analytic process and my data to my dissertation chair, my second committee member (both of whom are co-PIs on the seed grant that funded the work and both of whom have been present in the setting as guest lecturers), and my full committee, so that interpretive choices

that depend on my position can be tested against readers who share knowledge of the setting but do not share my investment.

This is the trustworthiness posture analytic autoethnography asks for: insider position acknowledged, multi-source evidence assembled, chair and committee engaged as interlocutors. I make no claim to neutrality, and analytic autoethnography does not require it.

B.7 Reflexivity performed

In this section I perform reflexivity rather than declare it. Three short reflexive memos make explicit how my position shaped my analytic choices, and a closing memo names my investment in the success of the work.

B.7.1 The K-12 to undergrad arc as researcher origin

My research began not in a university lab but in a high-school art classroom. I was an English-language-arts and high-school art teacher when I first encountered generative-AI tools as instructional resources. As part of that teaching practice I led a professional development for a school district in Colorado on how to incorporate DALL-E, Midjourney, and NightCafe in art and writing classrooms (RE-Q2). During my first year in the PhD program I designed and ran the Charles Burrell School of Arts diptych art contest, in which high-school students produced a hand-drawn image and a generative-AI-prompted image as paired pieces (RE-Q3, AP-2024-05-16). The contest drew more than sixty students.

This arc matters reflexively because it shaped what I noticed when I designed the CU undergraduate course. I came to the engineering classroom with a teacher's eye on accessibility, on student-produced work, on whole-child engagement, and on the rhetorical dimensions of prompting. I did not come from a computer-science research lab with a system-evaluation lens. The four-theme architecture that recurs across my iterations (Education, Industry, Ethics, Accessibility) is the visible trace of my teacher's eye. So is the recurrence of named student outputs in every opening session deck, the K-12 outreach work I sustained alongside the undergraduate teaching,

and the public-facing newsletter and podcast practice. Another researcher entering this work from a different origin would have noticed different things and made different curricular choices.

I name this so that a reader can calibrate how my position shaped the framework. The four-theme architecture is not a neutral analytic category. It is the architecture I built into my teaching from the K-12 origin onward, and I subsequently confirmed against external sources (the UW KidsTeam children's independent surfacing of the same themes in KT-THEMES, the Aspen Public Radio framing in AP-2024-05-16, the Luma feedback patterns in LF-3, and my online learners' responses in TR-4.D1 through TR-4.D5).

B.7.2 Navigating the practitioner-pioneer position

I am a practitioner-pioneer in the strict sense: I taught what I have documented to be the first generative-AI course at the University of Colorado Boulder (RE-Q3), and I sustained that pioneering position across the four iterations and the cross-iteration channels. This position carries its own reflexive demands.

First, the position is unrepresentative by design. Pioneer practice is rare, and no comparison-group instructor at CU Boulder was running the same course at the same time. My findings are therefore claims about pioneer practice rather than claims about typical instructor practice, and I frame them that way in Chapter D. The compression-as-curriculum-maturation claim, for example, is a claim about how pioneer practice matures under fast technological turnover, not a claim that any well-taught course of any topic must mature by compression.

Second, the position confers an information advantage. Because I have been the instructor at every site, I have access to materials no external researcher could obtain at the same depth (the Weekly Updates Prelim Document, the structured reflective journal, the public-facing newsletter, the podcast, the federal webinar). The advantage is real, and analytic autoethnography licenses its use. I make a point in §B.5.3 of subjecting each anchor concept to multi-source testing across artifact types so that the advantage does not become a license for unfalsifiable claims.

Third, the position interacts with my membership in the engineering-education program

(ENED) at CU Boulder. My disciplinary home is not computer science, and my course is taught from the College of Engineering side rather than from the CS curriculum. This positions my work within engineering education and HCI rather than within computer-science curriculum scholarship. The contributions I claim are calibrated to those literatures.

B.7.3 The methodological supplementation itself as autoethnographic data

The choice to supplement the main document's design-based-research analysis with an analytic-autoethnographic posture (§B.2) is itself an autoethnographic datum. The judgment to read my own iterations not only through DBR's curriculum-design lens but also through the autoethnographic lens of an instructor at the pioneer-entrant moment of a fast-moving field is a piece of practitioner-pioneer reasoning that the dissertation now treats as part of the work.

I narrate the supplementation here as part of the work, not as a backstage adjustment. The supplementation followed from my honest reading of what the four iterations were producing alongside the DBR analysis. The Weekly Updates Prelim Document recorded reflective observations (such as the Week 1 ChatGPT-quiz hallucination noting) that the DBR framing did not have a natural way to read as theoretical material. The cross-iteration channels (the Keep Up Newsletter, the podcast, the federal-research webinar) carried reflective content that sat outside DBR's curricular focus but was rich autoethnographic material. The choice to bring these into an autoethnographic supplement is itself the kind of move a complete-member researcher is positioned to make.

The supplementation also adapts the supplementary thematic-coding work the proposal had specified. The proposal anticipated a thematic-coding analysis of student-interview data that the iterations did not generate in IRB-governed form. Rather than retrofit thematic coding to the data the iterations did produce, this appendix substitutes a methodologically coherent alternative: analytic autoethnography reads the same artifact corpus that the DBR primary analysis reads but asks a different analytic question of it. The substitution is principled and is named here rather than concealed.

B.7.4 Acknowledgment of my investment in the work

I have a stake in this work succeeding. The dissertation is the document by which I become a doctor of philosophy in engineering education. The four-iteration course is the work I have built over three years. The findings I claim in Chapter D are claims I would like to land as scholarly contributions. Any reader of an analytic autoethnography must factor this investment in.

I respond to my investment in three ways. First, I source every substantive claim to specific artifacts and quote IDs so that any reader can audit the basis of the claim against the evidence-table document. Second, I name the asymmetries and limits of my data in §B.6.3 and §B.6.4 rather than concealing them. Third, I rely on my chair and my committee, who have visibility into the setting through their guest-lecture and seed-grant roles (and through my proposal-and-defense process), to test my interpretations against the same data I have. This is what analytic autoethnography asks of an invested practitioner-researcher.

I do not claim neutrality. I claim that my investment has been disciplined by multi-source evidence, by named limits, and by the chair-and-committee dialogue that has shaped this revised dissertation.

Appendix C

Iteration Narratives

C.1 Framing the iterations through analytic autoethnography

In this chapter I narrate the four iterations of my generative-AI course in turn. Each iteration narrative is sourced to the artifact corpus with named dates, named guests, named students, and named tools. The chapter closes on a cross-iteration comparative analysis (§C.6) that surfaces the patterns the iteration-by-iteration narration makes visible.

Figure C.1 below visualizes the practitioner-pioneer trajectory across six concurrent delivery channels for the entire period covered by this dissertation. The four iterations of the generative-AI course appear along the iteration spine (Iterations 1 and 2 as fifteen-week semester bars, Iterations 3 and 4 as compressed five-day workshops); the other channels (HCI graduate guest lectures, K-12 outreach, public-facing reflection through the Keep Up Newsletter and Podcast, and the CU RMACC federal-research webinar) appear in their own swimlanes. The figure orients the iteration narratives that follow.

C.1.1 What an analytic-autoethnographic narrative does with iteration data

Chapter B established the methodological framework I run on. Anderson's (2006) analytic autoethnography requires that the autoethnographer's insider position be the analytic resource, that reflexivity be analytic rather than confessional, and that the dialogue with informants beyond the self be substantive. These commitments shape how I narrate each iteration.

The narration is not a sequence of class-by-class summaries; it is a structured account of what

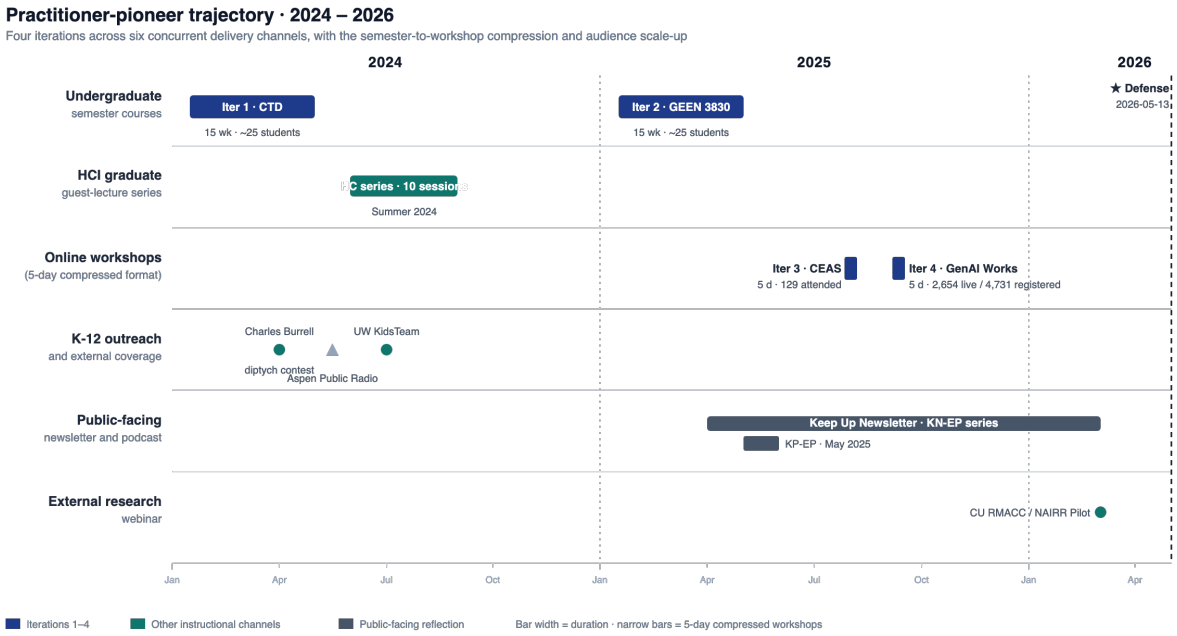

Figure C.1: Practitioner-pioneer trajectory across six concurrent delivery channels, January 2024 through May 2026. The four iterations of the generative-AI course (Iterations 1 and 2 as 15-week semester courses; Iterations 3 and 4 as 5-day compressed workshops) are visible alongside the HCI graduate guest-lecture series, K-12 outreach with Aspen Public Radio coverage and the UW KidsTeam collaboration, the Keep Up Newsletter and Podcast as public-facing reflection, and the CU RMACC federal-research webinar. Bar width is proportional to duration; narrow bars mark the five-day workshop compression of Iterations 3 and 4.

I built, what I observed, what I learned, and what I changed. The first-person voice is consistent throughout. Where I cite my own contemporaneous reflective journal (the Weekly Updates Prelim Document for Iteration 1, WU-1.W01 through WU-1.W15) I mark it as such. Where I cite my recorded teaching delivery (TR-4.D1 through TR-4.D5 for Iteration 4) I mark it as such. Where the iteration's reflective data is thinner (Iterations 2, 3, and 4), I name the limit and draw on the cross-iteration reflective channels (KN-EP series, KP-EP series, WB-2026-03-03) as the supplementary reflective base within an acknowledged retrospective-public frame.

C.1.2 Analytic posture as a precursor to the iteration narratives

The analytic posture I take across all four iterations is this. I treat each iteration as a connected component of a single multi-year practitioner-pioneer practice rather than as a discrete experimental trial. The unit of analysis is therefore the iteration-within-a-practice, not the iteration-as-trial. This matters because the design-conjecture-then-measure-then-refine logic of design-based research (which my proposal had specified and which I drew back from, §B.2) treats iterations as trials. The analytic-autoethnography logic I adopt treats iterations as moments in a continuing practitioner-pioneer trajectory.

The trajectory has shape. The first two iterations are semester-long undergraduate courses (CV-1 in Spring 2024, CV-2 in Spring 2025) that I treat as the consolidation phase of my pioneer practice. The second two iterations are five-day online workshops (DK-3 plus LF-3 and LR-3 in August 2025, DK-4 plus TR-4 series in September 2025) that I treat as the distillation phase. The compression-as-curriculum-maturation finding developed in §D.3 follows from this trajectory shape.

C.1.3 What each iteration narrative will cover

Each of the four iteration narratives in this chapter covers the same set of points:

- Institutional context (program, course code, audience size, sponsorship)
- Curriculum architecture as taught in that iteration
- Week-by-week or day-by-day structure where the artifacts allow

- Guest speakers and topics
- Named student work and learner outputs
- Tools introduced, retained, repositioned, or dropped
- Reflective material specific to the iteration
- Cross-iteration relationships (what the iteration learned from the previous and what it set up for the next)

The narrative depth varies by iteration depending on what the corpus supports. Iteration 1 has the most contemporaneous reflective journaling (WU-1 series); Iteration 4 has the most recorded teaching delivery (TR-4 series). I narrate each iteration with the depth its artifacts allow.

C.1.4 The patterns the iterations together exhibit

The iteration narratives in §C.2 through §C.5 surface patterns that Chapter D elaborates as autoethnographic sub-claims under the three curriculum-design principles named in the main dissertation document: modularity, learner choice, and continuous feedback. I preview the patterns here so the reader of the iteration narratives can hold the Chapter D destination in view while reading.

Architectural stability across changing surface · previews Modularity (§D.2). The four-theme architecture (Education, Industry, Ethics, Accessibility) appears in every iteration and across every delivery channel I documented. The tools change, the institutional sponsors change, the audience demographics change, the delivery duration compresses six-fold, but the architecture is stable. Chapter D §D.2 elaborates this as the inside view of the modularity principle.

Compression from semester to workshop · previews Modularity §D.2.3 (compression-as-curriculum-maturation). The shift from Iterations 1 and 2 (fifteen-week semester courses) to Iterations 3 and 4 (five-day online workshops) is the most consequential structural move in the practitioner-pioneer trajectory. The cross-iteration analysis in §C.6 documents the move in detail, and Chapter D §D.2.3 develops the compression-as-curriculum-maturation sub-claim as the

autoethnographic reading of that move under modularity.

Each iteration as a different choice-architecture for learners · previews Learner Choice (§D.3). Each iteration positions a different audience in a different relationship to choice over what to attend to, which tools to use, and how to demonstrate learning. Iteration 1 frames choice through a fifteen-week sequenced curriculum with a final project; Iteration 2 frames it through student teach-outs; Iteration 3 compresses choice into a five-day workshop with Luma feedback; Iteration 4 carries the same compression into a second cohort with day-by-day transcripts. Chapter D §D.3 elaborates this as the multi-channel teaching practice sub-claim under learner choice.

Continuous learner observation as the engine of revision · previews Continuous Feedback (§D.4). Across the iterations, what learners produced and said inside each iteration is what shaped the next. Weekly journal entries from Iteration 1 (WU-1 series), Iteration 3 Luma feedback (LF-3), and Iteration 4 day-by-day transcripts (TR-4 series) are the explicit records of this loop. Chapter D §D.4 elaborates this as the reflexive loop in practice and develops the hallucination-as-pedagogy sub-claim under continuous feedback.

Sections C.2 through 3.5 present the four iteration narratives in order. Section C.6 presents the cross-iteration comparative analysis.

C.2 Iteration 1 · CTD Pilot · Spring 2024

C.2.1 Institutional context

Iteration 1 was the first generative-AI course at the University of Colorado Boulder offered through the ATLAS Institute (RE-Q3). The course was hosted within the Creative Technology and Design (CTD) program, ATLAS's undergraduate program for radical creativity and innovation. The audience was approximately twenty-five CTD undergraduates. The duration was fifteen weeks: January 17, 2024 through May 1, 2024.

This was the discovery phase of my pioneer practice. I had no precedent at CU Boulder to

draw on; my prior teaching of generative-AI material had been in K-12 settings (RE-Q2 through RE-Q3) and in the Charles Burrell School of Arts diptych contest. The CTD pilot was where I built the four-theme curriculum architecture (§A.4) for the university-level setting.

C.2.2 Week 1 · January 17, 2024 · Prompt engineering opener

Week 1 opened with prompt engineering. The Week 1 deck (DK-1.W01) carries thirty-five slides and twenty-three named student prompt-engineering outputs from the opening session. The named outputs are the visible trace of the dialogue-with-informants-beyond-self criterion (§B.3.4) at the very beginning of the iteration.

Two of the twenty-three named students are traceable across the cross-iteration corpus. Ashley Stafford appeared in the Week 1 deck and was subsequently quoted in the Aspen Public Radio article (AP-2024-05-16) published one week after the iteration ended. Ethan Cuenca appeared in the Week 1 deck and was publicly credited in my Keep Up Newsletter Episode 3 (KN-EP3-Q1, "I learned about Soundful from one of my students during class") as the student whose tool discovery shaped Iteration 2's curriculum. The two cross-iteration cases are documented in §C.6.4 as named instances of the student-to-instructor tool flow that pioneer practice produces.

The Week 1 Weekly Updates entry (WU-1.W01) carries my contemporaneous reflective observations. The "Learned" section captures two distinct observations:

> WU-1.W01-Q1: "I learned that some of the multiple choice quizzes generated by ChatGPT were not correct and had hallucinations."

> WU-1.W01-Q2: "I learned that the students enjoyed creating the alter-ego images to describe themselves to the class."

The first observation is the seed of the hallucination-as-pedagogy finding I develop in §D.2. The second observation captures what the Midjourney Self-Portrait Assignment opened up: students using image generation to introduce themselves to their classmates as a low-stakes way to surface their own voices in the early weeks. Both observations were made in the first week of the first iteration.

C.2.3 Weeks 2 to 4 · Education theme

The Education theme occupied the first major curricular block of Iteration 1. Two guest speakers carried the theme:

- **Anthony Pinter** (ATLAS) on creativity and emerging-technology pedagogy
- **Nolan Brady and Shivendra** on Education applications of generative AI

The block also brought **Diane Sieber** (Herbst Program for Engineering, Ethics, and Society) into the course to discuss writing with generative AI. Sieber is a member of my dissertation committee and one of the two PIs on the AI-IRT Seed Grant that funded the broader research arc (AI-PROPOSAL). Her appearance as a guest in Iteration 1 was within the Education theme, framed around the writing question that her own teaching had been engaging.

C.2.4 Weeks 5 to 9 · Industry theme

The Industry theme carried generative-AI tools and practices into workplace contexts. Two guest speakers anchored the block:

- **Matt Zago** on video generation and industry applications of generative-AI video tools
- **Daniel Ritchie** of Hugging Face for a workshop session on open-model use

Ritchie's session was structured as a workshop rather than a lecture. The Canvas export (CV-1) records the session as a hands-on Hugging Face engagement that introduced students to model selection and use beyond the closed proprietary tools (DALL-E, ChatGPT, Midjourney) that the early weeks had emphasized.

C.2.5 Weeks 10 to 12 · Ethics block

The Ethics theme occupied a distinct block in Iteration 1. Two guest speakers carried it:

- **Nikolaus Klassen** (Google) on industry-internal ethics of generative AI

- **Suibi Weng** on Reality Editor and the question of mixed-reality augmentation

The Ethics block stands somewhat apart from the cross-cutting framing I describe in §A.4. The Iteration 1 curriculum gave Ethics a dedicated block as well as cross-cutting attention. In the Iteration 2 syllabus (SY-2) I committed more fully to the cross-cutting framing and removed the dedicated block. Section C.6.2 documents the structural shift between iterations.

C.2.6 Weeks 13 to 15 · Accessibility theme and Final Project

The closing block carried the Accessibility theme and the course's Final Project. The Final Project Requirements (FP-1) document a Media Studies-style company-creation assignment with nine reflection questions distributed across the four themes. The ethics question among them, "Do you think this will eliminate creative jobs?", made the cross-cutting Ethics theme present within the Accessibility-dominant final weeks.

The Final Project asked each student to design a company that used generative-AI tools and to explain the company's positioning within the four themes. The assignment produced student-level four-theme syntheses that became part of the dialogue-with-informants evidence base for §B.3.4.

Figure C.2 samples three of the curriculum materials preserved in the Iteration 1 Canvas export (CV-1). The materials span the semester: the Week 3 writing-with-GenAI assignment I authored, the Week 4 handout that Diane Sieber prepared for her guest discussion, and the closing Final Project Requirements that operationalized the four-theme architecture as the capstone assignment.

C.2.7 Aftermath · Aspen Public Radio coverage

One week after Iteration 1 ended, on May 16, 2024, Aspen Public Radio published an article titled "Could AI be the next college teaching assistant? Some Colorado professors believe so" (AP-2024-05-16). The article named me as the instructor of the course and quoted my Iteration 1 student Ashley Stafford on what the course had offered. The article's framing aligns with the Education

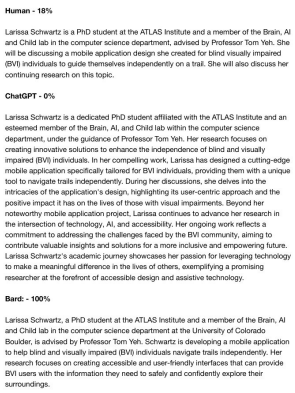

Human - 18%

Larissa Schwartz is a PhD student at the ATLAS Institute and a member of the Brain, AI and Child lab in the computer science department, advised by Professor Tom Yeh. She will be discussing a mobile application design she created for blind visually impaired (BVI) individuals to guide themselves independently on a trail. She will also discuss her continuing research on this topic.

ChatGPT - 0%

Larissa Schwartz is a dedicated PhD student affiliated with the ATLAS Institute and an esteemed member of the Brain, AI, and Child lab within the computer science department, under the guidance of Professor Tom Yeh. Her research focuses on creating innovative solutions to enhance the independence of blind and visually impaired (BVI) individuals. In her compelling work, Larissa has designed a cutting-edge mobile application specifically tailored for BVI individuals, providing them with a unique tool to navigate trails independently. During her discussions, she delves into the intricacies of the application's design, highlighting its user-centric approach and the positive impact it has on the lives of those with visual impairments. Beyond her noteworthy mobile application project, Larissa continues to advance her research in the intersection of technology, AI, and accessibility. Her ongoing work reflects a commitment to addressing the challenges faced by the BVI community, aiming to contribute valuable insights and solutions for a more inclusive and empowering future. Larissa Schwartz's academic journey showcases her passion for leveraging technology to make a meaningful difference in the lives of others, exemplifying a promising researcher at the forefront of accessible design and assistive technology.

Bard: - 100%

Larissa Schwartz, a PhD student at the ATLAS Institute and a member of the Brain, AI and Child lab in the computer science department at the University of Colorado Boulder, is advised by Professor Tom Yeh. Schwartz is developing a mobile application to help blind and visually impaired (BVI) individuals navigate trails independently. Her research focuses on creating accessible and user-friendly interfaces that can provide BVI users with the information they need to safely and confidently explore their surroundings.

(a) Week 3 writing-with-GenAI assignment (Larissa-authored, FP-precursor)

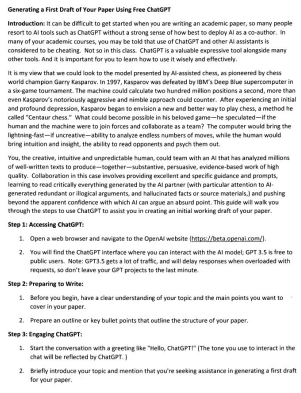

Generating a First Draft of Your Paper Using Free ChatGPT

Introduction: It can be difficult to get started when you are writing an academic paper, so many people resort to AI tools such as ChatGPT without a strong sense of how best to deploy AI as a co-author. In many of your academic courses, you may be told that use of ChatGPT and other AI assistants is considered to be cheating. Not so in this class. ChatGPT is a valuable expressive tool alongside many other tools. And it is important for you to learn how to use it wisely and effectively.

It is my view that we could look to the model presented by AI-assisted chess, as pioneered by chess world champion Garry Kasparov. In 1997, Kasparov was defeated by IBM's Deep Blue supercomputer in a six-game tournament. The machine could calculate two hundred million positions a second, more than even Kasparov's notoriously aggressive and nimble approach could counter. After experiencing an initial and profound depression, Kasparov began to envision a new and better way to play chess, a method he called "Centaur chess." What could become possible in his beloved game—he speculated—if the human and the machine were to join forces and collaborate as a team? The computer would bring the lightning-fast—if uncreative—ability to analyze endless numbers of moves, while the human would bring intuition and insight, the ability to read opponents and psych them out.

You, the creative, intuitive and unpredictable human, could team with an AI that has analyzed millions of well-written texts to produce—together—substantive, persuasive, evidence-based work of high quality. Collaboration in this case involves providing excellent and specific guidance and prompts, learning to read critically everything generated by the AI partner (with particular attention to AI-generated redundant or illogical arguments, and hallucinated facts or source materials,) and pushing beyond the apparent confidence with which AI can argue an absurd point. This guide will walk you through the steps to use ChatGPT to assist you in creating an initial working draft of your paper.

Step 1: Accessing ChatGPT:

1. Open a web browser and navigate to the OpenAI website (https://beta.openai.com/).
2. You will find the ChatGPT interface where you can interact with the AI model; GPT 3.5 is free to public users. Note: GPT3.5 gets a lot of traffic, and will delay responses when overloaded with requests, so don't leave your GPT projects to the last minute.

Step 2: Preparing to Write:

1. Before you begin, have a clear understanding of your topic and the main points you want to cover in your paper.
2. Prepare an outline or key bullet points that outline the structure of your paper.

Step 3: Engaging ChatGPT:

1. Start the conversation with a greeting like "Hello, ChatGPT!" (The tone you use to interact in the chat will be reflected by ChatGPT.)
2. Briefly introduce your topic and mention that you're seeking assistance in generating a first draft for your paper.

(b) Week 4 Writing with GenAI handout (Diane Sieber, guest)

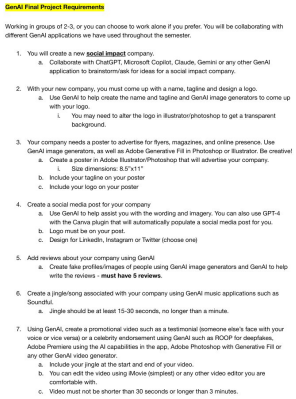

GenAI Final Project Requirements

Working in groups of 2-3, or you can choose to work alone if you prefer. You will be collaborating with different GenAI applications we have used throughout the semester.

1. You will create a new **social impact** company.
   a. Collaborate with ChatGPT, Microsoft Copilot, Claude, Gemini or any other GenAI application to brainstorm/ask for ideas for a social impact company.
2. With your new company, you must come up with a name, tagline and design a logo.
   a. Use GenAI to help create the name and tagline and GenAI image generators to come up with your logo.
      i. You may need to alter the logo in illustrator/photoshop to get a transparent background.
3. Your company needs a poster to advertise for flyers, magazines, and online presence. Use GenAI image generators, as well as Adobe Generative Fill in Photoshop or Illustrator. Be creative!
   a. Create a poster in Adobe Illustrator/Photoshop that will advertise your company.
      i. Size dimensions: 8.5"x11"
   b. Include your tagline on your poster
   c. Include your logo on your poster
4. Create a social media post for your company
   a. Use GenAI to help assist you with the wording and imagery. You can also use GPT-4 with the Canva plugin that will automatically populate a social media post for you.
   b. Logo must be on your post.
   c. Design for LinkedIn, Instagram or Twitter (choose one)
5. Add reviews about your company using GenAI
   a. Create fake profiles/images of people using GenAI image generators and GenAI to help write the reviews - **must have 5 reviews.**
6. Create a jingle/song associated with your company using GenAI music applications such as Soundful.
   a. Jingle should be at least 15-30 seconds, no longer than a minute.
7. Using GenAI, create a promotional video such as a testimonial (someone else's face with your voice or vice versa) or a celebrity endorsement using GenAI such as ROOP for deepfakes, Adobe Premiere using the AI capabilities in the app, Adobe Photoshop with Generative Fill or any other GenAI video generator.
   a. Include your jingle at the start and end of your video.
   b. You can edit the video using iMovie (simplest) or any other video editor you are comfortable with.
   c. Video must not be shorter than 30 seconds or longer than 3 minutes.

(c) Weeks 13–15 Final Project Requirements (FP-1)

Figure C.2: Three Iteration 1 curriculum materials preserved in the Canvas export (CV-1), spanning the semester: a Larissa-authored Week 3 writing assignment, the Week 4 guest handout from Diane Sieber, and the cross-iteration Final Project Requirements (FP-1).

theme of the four-theme architecture; its publication immediately after the iteration ended provides external corroboration that the iteration had landed pedagogically. I cite the article at source-level (AP-2024-05-16) in this dissertation.

C.2.8 What I learned across Iteration 1

The Weekly Updates Prelim Document (WU-1.W01 through WU-1.W15) carries my contemporaneous "Learned" entries across all fifteen weeks. The entries are short and structured: each week records what I had taught, what guests had come, what resources I had used, what I had learned, and what students had produced. The "Learned" entries are the most analytically valuable; they capture what the teaching had taught me, week by week, before retrospection set in.

Selected verbatim "Learned" entries that mark the arc of the iteration:

Week 2 (WU-1.W02-Q1, on tool accessibility and bias):

> "I learned that Midjourney is not free and can only be accessed through Discord. I learned that Midjourney images are far superior than any other GenAI image generators. I learned more about the bias associated with image generation within GenAI applications."

Week 3 (WU-1.W03-Q1, on AI content detection and institutional context):

> "I learned that AI content detectors are unreliable. I learned that every professor at the university has a different rule when it comes to using GenAI in their classroom."

Week 5 (WU-1.W05-Q1, on cross-industry application):

> "I also learned that various generative AI applications can benefit various industries, including surveillance, healthcare, marketing, advertising, education, gaming, media, podcasting, and more."

Week 6 (WU-1.W06-Q1, on the SORA video model and its limits):

> "I learned that SORA is going to be a game changer in the video/movie industry, but there are still a lot of issues that need to be addressed before releasing this product. I learned that GenAI is not only used creatively but can be used within all types of different industries, such as finance."

Week 7 (WU-1.W07-Q1, on deepfakes and student concern about creative originality):

> "I learned about Deepfakes and the ethical concerns surrounding them, as well as how to create them. I learned Deepfakes have been around for a while, as they used to be called Shallowfakes (using photoshop to alter images). I learned about the ethics surrounding music generation with GenAI. I learned that students are concerned with the originality of their music/art when using these different GenAI applications."

Week 14/15 (WU-1.W14-Q1, on accessibility and final-project execution):

> "I learned about different assistive technologies and applications that Derek uses daily, and which ones he recommends. I also learned about how he uses them for specific tasks. I learned how fast it took students to create marketing for a company from scratch, using different GenAI applications. I learned about new GenAI applications and the specific ones that the students preferred to use for their project."

Three patterns visible in the fifteen-week record are worth marking here.

First, hallucination as teachable phenomenon surfaced in the first week (WU-1.W01-Q1) and recurred throughout the semester as a teaching topic rather than as a problem to be hidden. The hallucination-as-pedagogy finding (§D.2) is anchored at this iteration's first week.

Second, the four-theme architecture stabilized through the iteration. The ForeverGold deck (DK-1.FG, slides 5-9) names the four themes as the course framing from the beginning; the iteration's week-by-week structure tested the framing against the curriculum as actually delivered. The architecture survived the iteration intact. The "Learned" entries above visibly trace the four themes: Week 3 covers Education (institutional policy on GenAI), Weeks 5 and 6 cover Industry (cross-industry application, SORA in the video industry), Week 7 covers Ethics (deepfakes and music ethics, student concern about originality), Weeks 14 and 15 cover Accessibility (Derek Riemer's lived experience with assistive technologies).

Third, the iteration produced student work that subsequently shaped Iteration 2. The Ethan Cuenca-to-Soundful flow (KN-EP3-Q1) is one named case; the iteration also produced student-to-tool discoveries that I carried into the Iteration 2 syllabus (SY-2) without always crediting

the student source in the curriculum itself. The flow is documented in §C.6.4 as the student-to-instructor tool flow pattern.

These three patterns establish Iteration 1 as the discovery phase of my pioneer practice. Iteration 2, narrated in §C.3, was the consolidation.

C.3 Iteration 2 · Mixed Engineering · Spring 2025

C.3.1 Institutional context

Iteration 2 ran from January 13, 2025 through April 30, 2025 as GEEN 3830-001 Special Topics in the College of Engineering and Applied Science (SY-2). The course was three credits, held in classroom DLC170, with an audience of approximately twenty-five undergraduates of mixed engineering backgrounds (computer science, mechanical, electrical, civil, and aerospace engineering, among others). The institutional move from ATLAS in Iteration 1 to the College of Engineering in Iteration 2 widened the student audience from the creativity-and-design CTD cohort to a broader engineering population.

This was the consolidation phase of my pioneer practice. I had spent a fifteen-week semester running Iteration 1 and had the contemporaneous reflective journal (WU-1 series) and the curricular artifacts (CV-1, DK-1.FG, DK-1.W01, FP-1) to draw on. Iteration 2 was the opportunity to test the four-theme architecture and the curriculum sequence against a new student population, with substantial tool turnover already required by the year of generative-AI change between Spring 2024 and Spring 2025.

C.3.2 Week 1 · January 13, 2025 · Prompt engineering opener (same as Iteration 1)

Week 1 of Iteration 2 opened with prompt engineering. The opening deck (DK-2.JAN13) preserved the structural commitment of Iteration 1: prompt engineering as the entry point, the Midjourney Self-Portrait Assignment as the Week 1 or Week 2 image-generation exercise, and

named student outputs as the visible trace of the dialogue-with-informants criterion.

The Midjourney Self-Portrait Assignment migrated forward to Week 1 itself in Iteration 2, slightly earlier than in Iteration 1 (where it was a Week 2 assignment). The Iteration 2 Midjourney deck (DK-2.JAN15) carries the assignment's framing for the engineering cohort.

The stability of the prompt-engineering opener across both iterations is one of the cross-iteration stable elements documented in §C.6.5.

C.3.3 Week 3 · DeepSeek added within weeks of public release

DeepSeek released its R1 model publicly in late January 2025. By Week 3 of Iteration 2 (early February 2025), I had incorporated DeepSeek into the curriculum. The Canvas export (CV-2) records the addition, and the syllabus (SY-2) lists DeepSeek among the tools students would engage. My dissertation chair Tom Yeh delivered a guest lecture on DeepSeek on January 28, 2025 (visible in CV-2 and on the Iteration 2 guest-lecturer spreadsheet AS-2.GUESTS).

DeepSeek's incorporation within weeks of its public release is one of the documented instances of immediate tool adoption that the cross-iteration analysis (§C.6.1) treats as evidence of how pioneer practice handles fast tool turnover.

Iteration 2 also added Claude 3.7 and Grok 3 as named tools in the Week 8 readings (per CV-2 module structure), bringing the major proprietary frontier-model line-up explicitly into the curriculum.

C.3.4 Weeks 7 to 9 · AI Agents as a full new module

Iteration 2 introduced a full module on AI Agents in Weeks 7 through 9, with the explicit assignment "Create your own AI Agent" as a hands-on component (per CV-2 module structure). The module did not exist in Iteration 1. The addition reflects how the field had matured between Spring 2024 and Spring 2025: AI Agents had become a workable pedagogical topic with concrete tool support and an addressable assignment scope, neither of which had been clearly possible a year earlier.

The AI Agents module is the single largest curricular addition between Iterations 1 and 2 and is one of the structural changes documented in §C.6.2.

C.3.5 Guest speakers as curated network in Iteration 2

The guest-speaker roster in Iteration 2 differed substantially from Iteration 1's, and the difference is itself part of the iteration's pedagogical contribution. Rather than treating guests as visiting one-offs, I curated the iteration's roster deliberately to track the changing technology landscape (DeepSeek and AI Agents were not topics anyone could have credibly taught in Spring 2024), to maintain CU-internal continuity (Nolan Brady recurred from Iteration 1 with a different topic), and to bring industry practice into the classroom (Justin Shacklette on GenAI for Software). The administrative spreadsheet AS-2.GUESTS records the iteration's guests; the four central appearances were:

Date	Guest	Topic	Curation rationale
2025-01-28	**Tom Yeh** (CU Boulder, Computer Science; my dissertation chair)	DeepSeek	New CU-internal guest brought in within weeks of DeepSeek's January 2025 public release; chose to teach the model from CU rather than rely on external industry-side coverage
2025-01-29	**Nolan Brady** (returning from Iteration 1)	GenAI in NeuroImaging	CU-internal continuity with a different topic; Brady's recurrence with a new lecture shows the curation is by-topic rather than by-person

Date	Guest	Topic	Curation rationale
2025-02-10	**Bobby Hodgkinson** (CU Boulder)	NotebookLM	New CU-internal guest covering the data-tooling and notebook-environment perspective
2025-02-12	**Justin Shacklette**	GenAI for Software	Industry guest bringing software-engineering practice into the classroom; the most direct industry-relevance move in the iteration

The curation work is the contribution at the network level, not just at the per-lecture level. Students did not encounter generative AI as an abstract category; they encountered it as a set of named practitioners and named topics actively assembled around the iteration's curriculum. The contrast with Iteration 1's roster — almost complete external-guest turnover, with CU-internal speakers stable but rotated to different topics — is documented in §C.6.3 as evidence of how the network is actively curated rather than passively maintained.

C.3.6 Student teach-out presentations

A pedagogical move new to Iteration 2 was the student teach-out: a class session in which a student took the front of the room and taught the class on a topic of their choosing. Six student teach-out presentations are preserved in the archive:

- **SP-2.DEEPFAKE** (Daniel Debretsion): DeepfakeAI as an Ethics-and-Industry topic
- **SP-2.SINTRA**: Sintra.ai as an Industry-and-Education topic, in Alt Style format
- **SP-2.HALEY** (Haley Phillips, two files)
- **SP-2.DAKOTA** (Dakota A)

- **SP-2.ROBOTICS**: The Integration of AI and Robotics as an Industry topic

The teach-out distribution across topics shows engineering-oriented choices weighted toward Industry-theme applications, with the DeepfakeAI presentation pulling Ethics-theme attention through the Industry surface. The teach-out artifacts are part of the dialogue-with-informants-beyond-self evidence base.

The administrative spreadsheets AS-2.TEACHOUT-DATES and AS-2.TEACHOUT-INDEX record the iteration's teach-out schedule and slide locations. The final-presentations schedule AS-2.FINAL-SCHED records the iteration's end-of-semester capstone presentations.

C.3.7 Structural reshuffles from Iteration 1

Iteration 2 was not Iteration 1 repeated. Several module-placement and curricular moves changed between iterations:

Element	Iteration 1	Iteration 2	Shift
Video module	Week 6	Week 14	Moved 8 weeks later
Industry theme block opening	Week 5	Week 8	Delayed 3 weeks
Hugging Face	Week 9	Week 7	Moved 2 weeks earlier
Reality Editor module (Suibi Weng)	Week 10	Not present	Dropped
AI Agents	Not present	Weeks 7 to 9	Added
Ethics block	Dedicated Weeks 10-12 plus cross-cutting	Cross-cutting only	Restructured

The reshuffles are documented in §C.6.2 as the module-placement-evolution pattern.

C.3.8 What I learned across Iteration 2

Iteration 2 had less structured contemporaneous reflective journaling than Iteration 1. The Keep Up Newsletter (KN-EP1, KN-EP2, KN-EP3, April through May 2025) and the Keep Up

Podcast (KP-EP2, KP-EP3, May 2025) carry my public-facing reflective writing across the Iteration 2 period, and analytic autoethnography permits this retrospective-public frame when it is named (§B.6.3). Three observations drawn from the iteration's artifacts and from these reflective channels are worth marking here.

First, the four-theme architecture survived the audience shift from CTD to mixed engineering. The Iteration 2 syllabus (SY-2) explicitly schedules the themes (Weeks 1-4 Education, Weeks 5-11 Industry, Weeks 12-15 Accessibility, with Ethics cross-cutting), and the iteration's curriculum delivered against the schedule. The architecture is not dependent on the creativity-and-design student type that Iteration 1's CTD cohort represented; it holds against an engineering-student type as well.

Second, the tool turnover pattern became visible. Adding DeepSeek within weeks of its January 2025 release (CV-2 records the Week 3 addition), bringing AI Agents in as a new module, adding Claude 3.7 and Grok 3, NotebookLM, Be My AI, Sora, HeyGen, ElevenLabs, and Custom GPTs with Wolfram API constituted a substantial tool refresh between iterations. The architecture absorbed the refresh without re-design of the themes.

Third, the cross-iteration guest-speaker turnover pattern became visible. External guests from Iteration 1 (Anthony Pinter, Diane Sieber, Matt Zago, Nikolaus Klassen, Suibi Weng, Daniel Ritchie) did not return as live guests in Iteration 2; CU-internal guests (Nolan Brady, joined now by Tom Yeh, Bobby Hodgkinson, and Justin Shacklette) carried the iteration. The pattern is documented in §C.6.3 and reflects how the instructor's actively-curated guest network adapts to the iteration's topic structure.

Iteration 2 completes the semester-length consolidation phase. The next two iterations (§C.4, §C.5) compressed the consolidated curriculum into five-day online workshops, opening the distillation phase.

C.4 Iteration 3 · GenAI in Five Cohort 1 · August 2025

C.4.1 Institutional context

Iteration 3 was the first compression of my course from the fifteen-week semester format to a five-day online workshop format. The workshop ran August 18 through August 22, 2025, at one hour of formal delivery per day, sponsored by the College of Engineering and Applied Science (CEAS) at CU Boulder. The workshop was hosted on Luma, a public event-registration and management platform, and was delivered through Zoom for the live sessions.

This was the first phase of distillation in my pioneer practice. With two semester-length iterations completed (CV-1 and CV-2), I had stabilized the four-theme curriculum architecture and developed a working sense of which content elements were essential to the pedagogical experience. Iteration 3 was the test of whether the architecture would survive a six-fold compression in delivery time.

C.4.2 Compression to five days at one hour per day

The compression from semester format (approximately forty-five hours of formal class time) to workshop format (five hours of formal delivery) is the most consequential structural move in my pioneer practice. The compression ratio is approximately nine-to-one in formal time and approximately six-to-one when informal time (chat engagement, asynchronous materials, recording replay) is included.

The compression-as-curriculum-maturation finding I develop in §D.3 follows from the move documented in this section.

C.4.3 Day-by-day structure · one topic per day

Iteration 3 used a one-topic-per-day structure across the five workshop days:

- **Monday** · Image generation
- **Tuesday** · Video generation

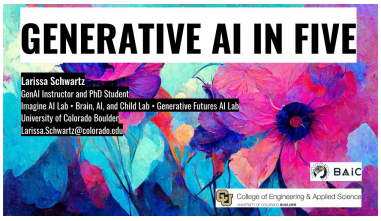

(a) Slide 1 (title)

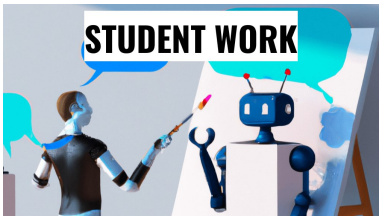

(b) Slide 4 (Monday)

(c) Slide 8 (Tuesday)

(d) Slide 12 (Wednesday)

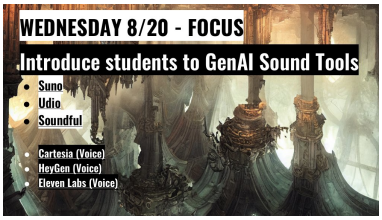

(e) Slide 16 (Thursday)

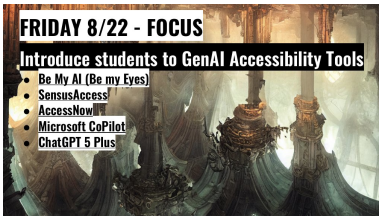

(f) Slide 20 (Friday)

Figure C.3: Selected slides from the Iteration 3 workshop deck (DK-3): title plus one representative slide from each day Monday through Friday.

- **Wednesday** · Audio generation
- **Thursday** · Research tools
- **Friday** · Human-centered AI

The workshop deck (DK-3) is twenty-two slides spread across the five days, with content density per slide higher than in the semester decks (DK-1.W01, DK-2.JAN13 through DK-2.MAR10). The compression operated at the slide level as well as at the duration level: more conceptual ground covered per slide, with the live demonstration carrying the explanatory load that lecture-discussion-lab cycles had carried in the semester iterations.

The day-by-day topic structure maps onto the four-theme architecture as follows: Monday and Tuesday sit primarily within the Industry-and-Education theme cluster (image and video tools as workplace and instructional applications). Wednesday sits within the Industry theme. Thursday sits within the Education theme (research tools as instructional aids). Friday delivers the Ethics-and-Accessibility synthesis as the Human-centered AI session that closes the workshop. The four-theme architecture is preserved through the day-by-day topical organization. Figure C.3 samples the deck across the five days, showing how the one-topic-per-day structure renders visually.

C.4.4 Audience composition · the Luma roster (LR-3)

The Iteration 3 Luma participant roster (LR-3) documents the audience composition with unusual specificity:

- **411 registered** total over the eight-month registration window
- **129 attended live** across the five days (matching the dissertation draft's "100+ live learners" claim)
- **65% students** of the registered population
- **70% expressed interest** in an AI Master's program in the registration form

The 411-to-129 attendance pattern shows substantial drop-off between registration and live attendance, which is typical for free online workshops. The 65% student figure indicates that the workshop reached primarily an undergraduate or graduate-student audience rather than an industry-professional one. The 70% Master's-interest figure indicates that the audience read the workshop as relevant to their formal AI education plans.

The audience demographic differs from the semester iterations' undergraduate-CU-Boulder population. Iteration 3 brought my curriculum to a broader online learner base while maintaining a student-heavy composition.

C.4.5 Learner feedback · the Luma survey (LF-3)

The Iteration 3 Luma feedback survey (LF-3) collected twenty-nine evaluative responses with the following distribution:

- 22 responses at five stars
- 6 responses at four stars
- 1 response at two stars
- Average rating: 4.69 of 5

Nine of the twenty-nine responses carried text feedback in addition to the numerical rating.

The text feedback is the principal qualitative learner data the iteration produced and is the dialogue-with-informants-beyond-self evidence base for this iteration (§B.3.4).

The nine text responses, verbatim, are:

ID	Rating	Verbatim text
LF-3-Q1	5★	"Best"
LF-3-Q2	5★	"I loved it!"
LF-3-Q3	5★	"I was wonderful session"
LF-3-Q4	5★	"It was really good"
LF-3-Q5	5★	"Excellent"
LF-3-Q6	5★	"Good many tools explored and learned now I'm trying one by one few tried"
LF-3-Q7	5★	"I really enjoyed the session by Larissa Schwartz, where she introduced us to the latest AI image generation tools like Midjourney, Microsoft Designer, DALL · E, Canva, Adobe Firefly, and NightCafe. She…" (response truncated in source)
LF-3-Q8	4★	"Very interesting."
LF-3-Q9	4★	"I think it was a good discussion regarding how to use the different AI image generation tools. A course work based on some of the neural networks behind them could be a great one."

Three things in this text-feedback corpus are worth marking for the analysis.

First, the brevity of most responses (one to three words for six of the nine) is characteristic

of post-event satisfaction surveys, which surface satisfaction rather than substantive engagement. The headline rating (4.69 of 5) is supported by these brief responses, and I treat them as one kind of evidence.

Second, the longer responses (LF-3-Q6, LF-3-Q7, LF-3-Q9) carry more analytical content. LF-3-Q6 names the breadth orientation as the workshop's strength ("many tools explored"). LF-3-Q7 names me by full name and lists the specific tools demonstrated in the image-generation day (Midjourney, Microsoft Designer, DALL · E, Canva, Adobe Firefly, NightCafe), confirming that the curriculum's tool catalog registered with at least one attentive learner. LF-3-Q9 is the constructive critique that surfaces the depth-versus-breadth tension developed in §C.4.6.

Third, the named-instructor recognition in LF-3-Q7 is itself a piece of analytic-autoethnographic evidence. The complete-member researcher (§B.3.1) is named by name by an external learner who attended the workshop and chose to identify me in their text feedback. This is dialogue with informants beyond self performed in the most literal way: the informant names the researcher.

C.4.6 The depth-versus-breadth tension surfaced by learner feedback

One four-star reviewer's verbatim text feedback surfaces a substantive curricular tension worth marking here:

> "I think it was a good discussion regarding how to use the different AI image generation tools. A course work based on some of the neural networks behind them could be a great one."

The response praises the workshop's tool-walkthrough orientation (the breadth move) while requesting a deeper engagement with the neural-network mechanics under the tools (the depth move). This is a depth-versus-breadth tension within the Education theme: the workshop format necessarily favors breadth (multiple tools demonstrated within one hour per day), and a reviewer engaged enough to leave constructive text feedback recognized what the format had foregone.

The tension is the kind of finding that analytic autoethnography surfaces from cross-context learner feedback. It is not an indictment of the workshop format; it is a recognition by an attentive

participant that the format makes a pedagogical trade-off, and a suggestion that future offerings could complement the breadth orientation with a depth-oriented sequel. The tension informs my future-directions discussion in §D.5.7.

C.4.7 What this iteration accomplished

Iteration 3 accomplished four things that set up Iteration 4 and supported the dissertation's findings.

First, it tested the four-theme architecture against the compression. The architecture survived. The Monday-through-Friday day-by-day topics map cleanly onto the theme clusters, and the workshop deck (DK-3) carries the architecture explicitly.

Second, it produced learner-facing data of a new kind. Where Iterations 1 and 2 produced student work (DK-1.W01 named outputs, SP-2 teach-outs, FP-1 final projects), Iteration 3 produced the Luma feedback corpus (LF-3) and the audience-composition data (LR-3). The data is structured and quantified in a way that the semester iterations' artifacts are not, and it provides triangulation for the curriculum claims that the semester iterations alone could not.

Third, it tested the pedagogical viability of the compressed format. The 4.69 of 5 average rating across twenty-nine responses indicates that the format was viable; the depth-versus-breadth tension surfaced in text feedback indicates the trade-off the format requires. Future-directions discussion in §D.5.7 addresses how the trade-off could be navigated in subsequent offerings.

Fourth, it set up Iteration 4 as a second-cohort delivery of essentially the same workshop. Iteration 4 (§C.5) used the same twenty-two-slide template (DK-4 is structurally identical to DK-3 with date headers changed) and the same Monday-through-Friday topical organization. The two workshop iterations are the distillation phase of my pioneer practice, with Iteration 3 establishing the format and Iteration 4 extending it to a global audience through the GenAI Works partnership.

Iteration 3 closed the consolidation-to-distillation transition. Iteration 4, narrated in §C.5, opened the multi-channel reach that §D.4 develops as the multi-channel teaching practice finding.

C.5 Iteration 4 · GenAI Works Cohort · September 2025

C.5.1 Institutional context

Iteration 4 was the second compressed cohort of my course, delivered September 8 through September 12, 2025, through partnership with GenAI Works. The workshop was delivered via the GenAI Works YouTube channel rather than through a closed institutional Zoom session, opening the audience to global online attendees during live broadcast. The structure preserved the Monday-through-Friday day-by-day format that Iteration 3 had established.

This iteration extended the distillation phase of my pioneer practice and is the iteration with the most extensive recorded teaching delivery in the corpus. The five day-by-day transcripts (TR-4.D1 through TR-4.D5) together total approximately fifty-five thousand words of verbatim teaching delivery.

C.5.2 The same template as Iteration 3 with date headers changed

The Iteration 4 workshop deck (DK-4) preserves the structure of the Iteration 3 deck (DK-3). The deck carries twenty-two slides with the same day-by-day organization (Monday image, Tuesday video, Wednesday audio, Thursday research, Friday human-centered AI), the same topical structure within each day, and only the date headers updated to reflect the September 8 through September 12 schedule.

The template stability between cohorts is the most concrete evidence of the compression-as-curriculum-maturation pattern (§D.3). The same workshop runs again, with the same architecture and the same daily topics, for a new audience three weeks after Cohort 1 concluded. The pedagogical structure had stabilized through Iteration 3's first delivery and operated as a working artifact in Iteration 4. Figure C.4 samples DK-4 across the same five days as the DK-3 gallery in §C.4.3; the two galleries should be read side-by-side to see how the template held its shape between cohorts.

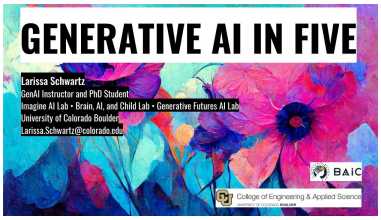

(a) Slide 1 (title)

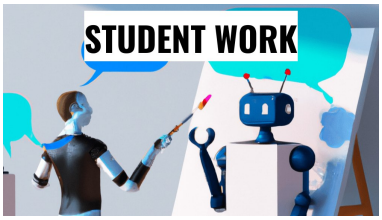

(b) Slide 4 (Monday)

(c) Slide 8 (Tuesday)

(d) Slide 12 (Wednesday)

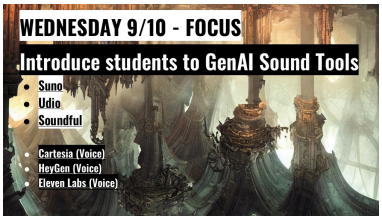

(e) Slide 16 (Thursday)

(f) Slide 20 (Friday)

Figure C.4: Selected slides from the Iteration 4 workshop deck (DK-4): same template as DK-3 with date headers updated to September 8 through 12, 2025.

C.5.3 YouTube delivery and global audience visibility

The GenAI Works YouTube delivery channel made the workshop visible to a globally-distributed audience during live broadcast. We originally configured the workshop on Streamyard, but the platform could not handle the registered audience capacity, and we switched to streaming the live broadcast on YouTube Live to reach the live audience at scale. The audience data exceed the Cohort 1 figures substantially: the Iteration 4 Luma event page accumulated approximately 4,731 registered guests (LR-4) across the cohort's promotion window, and the Day 1 YouTube broadcast on September 8, 2025 drew 2,654 live online participants (LR-4.D1). Audience feedback collected through the Luma platform produced 256 ratings with an average of 4.2 of 5 (LF-4), a higher ratings volume than Cohort 1's 29-response Luma feedback corpus (LF-3), reflecting the substantially larger registered base.

The Day 1 transcript (TR-4.D1, 10,962 words) opens with the host Samuel Cummings welcoming participants and naming the geographies in the chat (TR-4.D1-Q1, verbatim):

> "Amazing. Welcome everybody. I mean, I see so much excitement in the chat. How's everybody doing out there? … There we go. We got um people from all over the world. I see Nigeria, Denver, UK, Costa Rica. I mean, what a time to be alive, to be here virtually."

The opening also names me by name and introduces my dissertation chair Tom Yeh as a co-host from CU Boulder (TR-4.D1-Q2, verbatim):

> "So, uh we have Tom Yei. welcome from uh you know University of Colorado uh Boulder. We have Lissa Schwarz as well is going to be our main host today."

(Speech-recognition errors in the transcript include "Tom Yei" for Tom Yeh and "Lissa Schwarz" for Larissa Schwartz; the names are correctly identifiable in context.)

The audience identification is performed live in the workshop format because GenAI Works prompts each attendee to identify themselves in the chat. This is a piece of evidence for the multi-channel teaching practice finding (§D.4): the same four-theme curriculum that had reached

approximately twenty-five Boulder CTD undergraduates in Iteration 1 now reached a globally-distributed online audience in Iteration 4, with the architecture preserved.

C.5.4 Day-by-day transcript analysis

The five day-by-day transcripts are the principal artifact of Iteration 4. Their volumes are:

ID	Day	Topic	Approximate word count
TR-4.D1	Monday	Image generation	10,962
TR-4.D2	Tuesday	Video generation	11,395
TR-4.D3	Wednesday	Audio generation	11,599
TR-4.D4	Thursday	Research tools	9,921
TR-4.D5	Friday	Human-centered AI and Vibe Coding	11,276

Total: approximately 55,000 words of verbatim teaching delivery across the five days.

The transcripts capture my live instructional voice in ways that the slide decks and the Luma feedback survey for Iteration 3 do not. They are the principal evidence base for the narrative-visibility criterion (§B.3.3) within Iteration 4 and the contemporaneous record of how the compressed workshop format actually unfolded in delivery. The transcripts carry evidence for each of the three curriculum-design principles that Chapter D elaborates.

Modularity in the day-level structure (previews §D.2). Each day of Iteration 4 corresponds to a distinct curriculum module: image generation (D1), video generation (D2), audio generation (D3), research tools (D4), and human-centered AI with Vibe Coding (D5). The day-level packaging is the same modular architecture I had run across the four-theme weekly arc in Iterations 1 and 2, compressed into a five-day delivery. The transcripts let the modular structure be read as instructor-voice delivery rather than as static slide framing. Chapter D §D.2 elaborates the architectural-stability reading of this evidence; §D.2.3 elaborates the compression-as-curriculum-maturation sub-claim that the five-day cadence supports.

Learner choice instantiated tool-by-tool (previews §D.3). Across the 55,000 words I

walk learners through multiple tools per topic, invite the audience to choose among them, and treat the choice itself as a learning move. The Microsoft Designer walkthrough on Day 1 (TR-4.D1-Q3) closes on the choice frame: "if you're still not getting a result that you like, then try a different tool." Chapter D §D.3 elaborates the dialogue-with-informants reading of this evidence, and §D.3.3 elaborates the multi-channel teaching practice sub-claim that the global-audience delivery surface supports.

Continuous feedback as real-time response to encounters (previews §D.4). The transcripts document the reflexive loop in action: I observe what a tool produces, name what I see, and adjust the pedagogy in the next sentence rather than in the next iteration. The TR-4.D1-Q3 passage is one explicit instance (verbatim):

> "And you can also see these different hallucinations that are going on. So there aren't even bodies in these shoes. And so it's really interesting how these different tools are used. Sometimes I get some really really great image generation and it just depends on what you want to use it for, but other times there are a lot of hallucinations that you can get within these specific tools. And so once again, you just have to make sure and take a look at everything. And even if you do get a hallucination on the first time, it doesn't mean give up. You can keep inputting and inputting and inputting the prompts into into the generator and see if it changes. And so I would I would continue to keep working with the specific tool and if you're still not getting a result that you like, then try a different tool."

I name the phenomenon by its technical term (hallucination), point to a concrete example (the shoes without bodies), normalize the occurrence, and convert the encounter into an iteration practice in real time. Chapter D §D.4 elaborates this passage as evidence for the continuous-feedback principle, with hallucination-as-pedagogy treated as a side sub-claim that surfaces inside the principle rather than as the principle-level finding. The principle the transcript supports is continuous feedback; hallucination-as-pedagogy is one place where that principle becomes visible in instructor practice.

The transcripts are cited at source-level (TR-4.D1 through TR-4.D5) throughout this dissertation, with Q-IDs assigned for specific instructor-voice passages. The master evidence-table records the full set of Q-IDs in use.

C.5.5 Tom Yeh as guest from CU Boulder on Day 1

Day 1 of Iteration 4 featured my dissertation chair Tom Yeh as a guest speaker connecting from CU Boulder (visible in TR-4.D1). Yeh's appearance is one of the cross-iteration member-checking moments noted in §B.5.4: he is one of the two PIs on the AI-IRT Seed Grant (AI-PROPOSAL) that funded the broader research arc, and his presence as a guest in Iteration 2 (DeepSeek lecture, January 28, 2025) and now Iteration 4 (Day 1 guest) is partially documented in the corpus itself.

The CU-Boulder-to-global-audience pattern that Yeh's guest appearance instantiates is also a piece of evidence for the multi-channel teaching practice finding. The local institution (CU Boulder) and the global delivery channel (GenAI Works on YouTube) were brought into the same live session through the guest configuration.

C.5.6 Comparison to Cohort 1 · same structure, different delivery surface

Iteration 4 and Iteration 3 are structurally identical workshops delivered to two different audiences through two different platforms. The differences worth marking are:

Dimension	Iteration 3	Iteration 4
Sponsoring institution	CEAS (CU Boulder College of Engineering and Applied Science)	GenAI Works (external partnership)
Delivery platform	Luma plus Zoom	GenAI Works YouTube Live (Streamyard was tried first but could not handle the registered audience capacity)
Audience composition	65% students, primarily Boulder-and-US	Global, with Nigeria, UK, Costa Rica, and US attendees named in Day 1 transcript opening

Dimension	Iteration 3	Iteration 4
Live audience size	129 attended (LR-3)	2,654 live Day 1 participants on YouTube Live (LR-4.D1), against 4,731 Luma registrations (LR-4)
Audience evaluation	Luma feedback survey, 29 responses, 4.69 of 5 average (LF-3)	Luma feedback survey, 256 ratings, 4.2 of 5 average (LF-4)
Recorded delivery available?	Limited	Comprehensive (five day-by-day transcripts, 55,000 words)
Guest speakers	Curriculum-internal only	Tom Yeh from CU Boulder on Day 1

The differences are substantial in audience and platform; the curricular structure is the same. The same workshop runs against two different audience-and-platform configurations.

C.5.7 What this iteration accomplished

Iteration 4 accomplished four things that complete the dissertation's evidence base.

First, it extended the distillation phase to a global audience. The compression-as-curriculum-maturation pattern (§D.3) is documented across both Iteration 3 (CEAS-sponsored, predominantly student audience) and Iteration 4 (GenAI Works partnership, global audience). The architecture survives both deliveries.

Second, it produced the corpus's most extensive recorded teaching delivery. The five day-by-day transcripts (TR-4.D1 through TR-4.D5) total approximately fifty-five thousand words, providing the analytic-reflexivity criterion (§B.3.2) with a level of contemporaneous delivery evidence that Iterations 1, 2, and 3 do not match individually.

Third, it provided the most recent anchor source for the hallucination-as-pedagogy finding (§D.2). The Day 1 transcript (TR-4.D1) carries hallucination as a teaching topic in a stabilized form, completing the finding's twenty-month evidence base across four independent sources.

Fourth, it instantiated the multi-channel teaching practice finding (§D.4) at scale. The pioneer's four-theme curriculum reached a global online audience through a non-CU delivery platform, with the dissertation chair appearing as a guest from the home institution. The multi-channel character of the pioneer practice is most fully visible in this iteration.

Section C.6 next presents the cross-iteration comparative analysis that synthesizes the four iteration narratives.

C.6 Cross-iteration comparative analysis

In this section I step back from the iteration-by-iteration narratives (§C.2 through §C.5) and surface the patterns that the four iterations together exhibit. The patterns are five: tool turnover, module-placement evolution, guest-speaker turnover, student-to-instructor tool flow, and stable elements that persist across the iterations.

C.6.1 Tool turnover across iterations

The tool list shifted substantially between iterations, driven by the rapid public release of new generative-AI tools across 2024 and 2025. Comparing the Canvas LMS exports (CV-1 and CV-2) and the workshop decks (DK-3 and DK-4) surfaces specific tools added, repositioned, and dropped.

Tools added in Iteration 2 (Spring 2025) that were not present in Iteration 1 (Spring 2024):

- **DeepSeek** · added Week 3 within weeks of its January 2025 public release
- **Claude 3.7** · added to Week 8 readings
- **Grok 3** · added to Week 8 readings
- **AI Agents** · introduced as a full new module across Weeks 7 to 9, with the explicit assignment "Create your own AI Agent"
- **NotebookLM** · moved from student-mentioned in Iteration 1 to dedicated guest lecture

(Bobby Hodgkinson, February 10, 2025) in Iteration 2

- **Be My AI** · added to the Accessibility-theme block as a lived-experience-relevant accessibility tool
- **Sora** · added to the video-generation block as the OpenAI video model
- **HeyGen** · added to the video-generation block as the avatar-and-presentation tool
- **ElevenLabs** · added to the audio-generation block as the voice-cloning and text-to-speech tool
- **Custom GPTs with Wolfram API** · introduced as a workflow that combined Custom GPTs with the Wolfram computational backend

Tools repositioned between iterations:

- **Hugging Face** · moved from Iteration 1 Week 9 (Daniel Ritchie's live workshop) to Iteration 2 Week 7 (with Ritchie's video assigned for the iteration rather than a live appearance)

Tools dropped between iterations:

- **Reality Editor** · present in Iteration 1 Week 10 with Suibi Weng as live guest; absent from Iteration 2

Tools that recurred across all four iterations:

- **ChatGPT** (and successors in the OpenAI line)
- **Midjourney** (used in the Self-Portrait Assignment across all four iterations)
- **DALL-E** (used at the STEAM Festival mural and across the iterations)
- **Soundful** (introduced via Iteration 1 student Ethan Cuenca, retained through Iteration 4)
- **Prompt engineering** (as a meta-tool, anchoring Week 1 of every iteration)

The pattern is substantial tool turnover (more than ten new tools introduced in Iteration 2 alone) absorbed by a stable four-theme architecture. The architecture is the part that does not change; the tool list is the part that does.

C.6.2 Module-placement evolution

Beyond tool turnover, module placement within the semester structure shifted between Iterations 1 and 2:

Module	Iteration 1 placement	Iteration 2 placement	Shift
Prompt engineering opener	Week 1	Week 1	Stable
Midjourney Self-Portrait	Week 2	Week 1 (DK-2.JAN15)	Earlier
Education theme block	Weeks 2-4	Weeks 1-4	Stable to slightly earlier
Industry theme block opening	Week 5	Week 8	Delayed 3 weeks
Hugging Face	Week 9 (Ritchie live)	Week 7 (Ritchie video)	2 weeks earlier; live to recorded
Video module	Week 6	Week 14	8 weeks later
Ethics dedicated block	Weeks 10-12 (Klassen, Weng)	Cross-cutting only	Restructured to cross-cutting
Reality Editor	Week 10 (Weng)	Absent	Dropped
AI Agents module	Absent	Weeks 7-9	Added
Accessibility theme block	Weeks 13-15	Weeks 12-15	Slightly earlier
Final Project	Weeks 13-15	Final weeks	Stable

The video module's eight-week shift (Iteration 1 Week 6 to Iteration 2 Week 14) is the largest single structural move. Several factors plausibly informed the move (the maturity of Sora and HeyGen as classroom-usable tools in 2025 that had not been classroom-usable in 2024; the placement of video as the closing technical theme before the Final Project), and my retrospective reflective memo for Iteration 2 develops the reasoning.

The compression-and-stabilization pattern documented in §D.3 emerges from the table. Most modules are stable in placement; a small number shift substantially; and the shifts cluster around modules whose tool support changed materially between iterations.

C.6.3 Guest-speaker network curation as an instructional move

Across the two semester iterations I assembled and revised a network of guest practitioners that together carried the curriculum's industry, ethics, and accessibility connections into the classroom. The curation pattern is itself a finding about how this pioneering practice operates. The roster is not a fixed lecture series; it is an actively curated rotating network, and the curation work is part of what the instruction consists of.

Iteration 1 guests	Iteration 2 guests
Anthony Pinter (ATLAS, creativity)	(External, not returning live)
Diane Sieber (Herbst, writing)	(External, not returning live)
Matt Zago (finance-industry video)	(External, not returning live)
Daniel Ritchie (Hugging Face workshop)	Daniel Ritchie's video assigned to Week 7 (not live)
Nikolaus Klassen (Google ethics)	(External, not returning live)
Suibi Weng (Reality Editor)	(External, not returning)
Nolan Brady plus Shivendra (Education)	Nolan Brady on GenAI in NeuroImaging (returning CU guest with different topic)
—	Tom Yeh on DeepSeek (new CU-internal guest)
—	Bobby Hodgkinson on NotebookLM (new CU-internal guest)
—	Justin Shacklette on GenAI for Software (new industry guest)

Three properties of the curation are worth naming.

1. Industry relevance was sustained across iterations through deliberate rotation. Iteration 1 carried Matt Zago on finance-industry generative-AI applications and Nikolaus Klassen on industry-internal ethics from Google; Iteration 2 carried Justin Shacklette on software engineering and Bobby Hodgkinson on notebook environments. The specific industries shifted (finance, ethics, software engineering, data tooling) but the industry-presence commitment was preserved. Students at every iteration encountered generative AI as a practice with industry stakes and named

industry practitioners, not as a research curiosity.

2. CU-internal continuity sustained the network across topic turnover. Nolan Brady's recurrence from Iteration 1 to Iteration 2 with a different lecture topic shows that the network's continuity is by-relationship, not by-fixed-content. Tom Yeh's appearance in Iteration 2 (DeepSeek) and again in Iteration 4 (Day 1 of the GenAI Works workshop, from CU Boulder) extends the same pattern across the iterations. The network's CU-internal core can be re-aimed at new topics as the technology landscape moves.

3. External speakers cycle out at high rates as topics evolve. Iteration 1's external speakers (Pinter, Sieber, Zago, Klassen, Weng) did not return live in Iteration 2. This is not a failure of relationship maintenance; it is a feature of the curation. When the technology landscape shifts (DeepSeek and AI Agents emerge between Spring 2024 and Spring 2025), the appropriate external voices shift with it. A fixed external roster would have meant covering 2024 topics in a 2025 course.

The curation work is therefore part of the pedagogical contribution, not infrastructure separate from it. Documenting **who I brought in, when, and why** is documenting how a pioneer instructor sustains industry connection and topical currency across iterations. The administrative spreadsheets (AS-2.GUESTS, AS-2.TEACHOUT-DATES, AS-2.TEACHOUT-INDEX, AS-2.FINAL-SCHED) are the operational trace of this curation work in Iteration 2; Iteration 1's curation lives in the Canvas export (CV-1) and the Weekly Updates Prelim Document.

C.6.4 Student-to-instructor tool flow · the Ethan Cuenca to Soundful case

One of the most analytically interesting cross-iteration patterns is the documented flow of a tool from a student in Iteration 1 into the instructor's curriculum for Iterations 2, 3, and 4. The flow is documented through three independent artifacts:

1. **Slide 13 of the Iteration 1 Week 1 deck (DK-1.W01)** names Ethan Cuenca as one of the twenty-three students whose prompt-engineering outputs are displayed in the opening

session.

2. **The Weekly Updates Prelim Document (WU-1 series)** notes that one student presented on the GenAI music tool Soundful, and that I incorporated the student's discovery into the class assignment.

3. **Keep Up Newsletter Episode 3 (KN-EP3-Q1)** publicly credits the flow:

   "I learned about Soundful from one of my students during class."

Following Iteration 1, Soundful appeared in the Iteration 2 curriculum (visible in CV-2), in the Iteration 3 workshop deck (DK-3), and in the Iteration 4 workshop delivery (TR-4.D3 carries Soundful as part of the Audio-generation day). The tool persisted across all four iterations following its student-discovery introduction.

This is the autoethnographic claim "student co-discovery shaped the curriculum" made concrete at the level of a single named instance: a named student, an Iteration 1 dated artifact, a contemporaneous reflective note, a public-facing acknowledgment, and persistent presence in the subsequent curriculum. I want to be precise about what this is and is not: it is one well-documented case, not a pattern with multiple documented instances. Other student tool discoveries shaped the curriculum across the iterations (the Iteration 1 Final Project, the Iteration 2 teach-out presentations, the Iteration 3 Luma feedback all surfaced tool suggestions), but they are not documented with the same precision as the Cuenca-to-Soundful flow. The single named case speaks to the dialogue-with-informants-beyond-self criterion (§B.3.4) at the cross-iteration scale; it is offered as a concrete instance of a broader phenomenon I observed but did not document case-by-case. Figure C.5 samples DK-1.W01 to show the named-student-output slides on which the Cuenca case rests.

C.6.5 Stable elements across all four iterations

What does not change across the four iterations is as analytically important as what does. Five elements are stable from Iteration 1 through Iteration 4:

(a) Slide 1 (title)

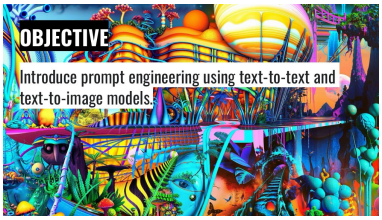

(b) Slide 5

(c) Slide 13 (Ethan Cuenca)

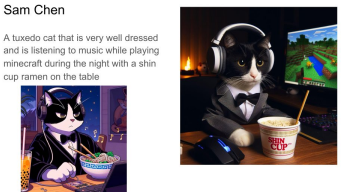

(d) Slide 18

(e) Slide 25

(f) Slide 31

Figure C.5: Selected slides from the Iteration 1 Week 1 opening lecture (DK-1.W01, January 17, 2024). Among the 23 named student prompt-engineering outputs displayed in the session, slide 13 features Ethan Cuenca, whose Soundful discovery shaped the curriculum across Iterations 2 through 4 and was publicly credited in KN-EP3-Q1.

The four-theme architecture. Education, Industry, Ethics, Accessibility, with Ethics increasingly positioned as cross-cutting across the others. The architecture is named explicitly in DK-1.FG (slides 5-9) for Iteration 1, in SY-2 for Iteration 2, in DK-3 for Iteration 3, and in DK-4 for Iteration 4.

The Week 1 prompt-engineering opener. Iteration 1 (DK-1.W01), Iteration 2 (DK-2.JAN13), Iteration 3 (DK-3 Day 1), and Iteration 4 (DK-4 Day 1, TR-4.D1) all open with prompt engineering. The opener is the most stable single curricular element in the corpus.

The Midjourney Self-Portrait Assignment. Used in Iteration 1 (Week 2), Iteration 2 (Week 1, DK-2.JAN15), and the workshops (Day 1 of Iterations 3 and 4 within the image-generation block). The assignment is the most stable single hands-on element.

Hallucination as teachable topic. Appears in Iteration 1 Week 1 reflection (WU-1.W01-Q1), in cross-iteration public writing (KN-EP1-Q1), and in the most recent workshop delivery (TR-4.D1). The hallucination-as-pedagogy finding (§D.2) is anchored at this stable element.

Named student outputs as opening-session content. The dialogue-with-informants criterion is performed at the opening of every iteration through named student work. Iteration 1's twenty-three named outputs in DK-1.W01 set the pattern; subsequent iterations continued the practice of naming student contributions in opening-session content.

The stable elements together constitute the durable pedagogical contribution of the four iterations. The architecture survived, the Week 1 opener survived, the Self-Portrait Assignment survived, the hallucination treatment survived, and the practice of naming student contributions survived. These are the elements that the cross-iteration distillation kept.

C.6.6 Compression in the timeline

The final cross-iteration pattern is the compression of delivery time:

Iteration	Duration	Formal class time
1 · CTD Pilot	15 weeks	~45 hours

Iteration	Duration	Formal class time
2 · Mixed Engineering (GEEN 3830)	15 weeks	~45 hours
3 · GenAI in Five Cohort 1	5 days	~5 hours
4 · GenAI Works Cohort	5 days	~5 hours

The compression ratio from semester to workshop is approximately nine-to-one in formal class time. The compression preserved the four-theme architecture, the Week 1 prompt-engineering opener, the Midjourney Self-Portrait Assignment, the hallucination-as-pedagogy treatment, and the practice of dialogue with informants beyond self. The compression-as-curriculum-maturation finding (§D.3) is the analytic claim that this compression represents.

C.6.7 Synthesis

The four iterations exhibit a coherent practitioner-pioneer trajectory. Iteration 1 was the discovery phase. Iteration 2 was the consolidation phase. Iterations 3 and 4 were the distillation phase. Across the trajectory, the four-theme architecture stabilized as the durable contribution, the tools turned over substantially, the guest network adapted, the student-to-instructor tool flow operated as a named case, and the delivery compressed approximately six-fold without content loss.

Chapter D develops the analytic claims that emerge from this trajectory: hallucination-as-pedagogy (§D.2), compression-as-curriculum-maturation (§D.3), and multi-channel teaching practice (§D.4).

Appendix D

Autoethnographic Findings

D.1 Framing through autoethnographic elaboration of the three principles

This chapter applies the supplementary autoethnographic posture (established in Chapter B) to the three curriculum-design principles the main document develops: **modularity**, **learner choice**, and **continuous feedback**. The chapter does not introduce three competing findings. It elaborates on each of the three principles by drawing on the instructor's reflexive position to surface what the DBR analysis of the main document does not. The autoethnographic supplement deepens; it does not replace.

D.1.1 The relationship between the three principles and the autoethnographic elaborations

The main document presents the three curriculum-design principles as the iterative-refinement work's primary outputs. Each principle is supported by evidence across the iterations; each is actionable for adopters; each is empirical in the DBR sense. What the principles do not by themselves convey is what the instructor's reflexive position on her own practice surfaces when she reads the same iterations autoethnographically.

This chapter provides that reflexive layer for each principle in turn. The three principles and their autoethnographic elaborations are:

Principle (main document)	Autoethnographic elaboration (this chapter)
Modularity	The architecture's stability seen from inside: the practitioner-pioneer's record of **what survives** when the surface changes (§D.2). Sub-claim: compression-as-curriculum-maturation.
Learner Choice	Dialogue with informants across channels: the practitioner-pioneer's record of **how learners shape the curriculum** across multiple delivery contexts (§D.3). Sub-claim: multi-channel teaching practice.
Continuous Feedback	The reflexive loop in practice: the practitioner-pioneer's record of **how moment-to-moment encounters with tools and learners shape revision** (§D.4). Sub-claim: hallucination-as-pedagogy.

Each elaboration carries its own sub-claim (named in the right column). The sub-claims are not separate findings competing with the principles; they are concrete patterns that the autoethnographic analysis surfaces within each principle.

D.1.2 Anderson's theoretical analysis criterion as the lens

Anderson's (2006) fifth criterion for analytic autoethnography is commitment to theoretical analysis. The autoethnographer's insider position is not used to produce evocative description; it is used to generate, refine, or extend theoretical understanding of broader social phenomena. This chapter applies the criterion to each of the three principles, drawing theoretical depth from the practitioner's reflexive vantage that the DBR analysis of the main document does not surface.

What the criterion governs is how each elaboration counts as theoretical analysis. An elaboration must name a generalizable pattern that the insider position made possible to see, support that pattern against multiple independent sources, and have implications beyond the four iterations

for how scholarship treats the phenomenon.

The three sub-claims meet these tests. **Compression-as-curriculum-maturation** is a pattern about how the modular architecture matured by distillation rather than by accretion; the corpus supports it across the four iterations. **Multi-channel teaching practice** is a pattern about how learner choice operated across eight contexts of varying duration, audience, and depth; the corpus documents it across more than two and a half years. **Hallucination-as-pedagogy** is a pattern about how moment-to-moment continuous feedback with the tools became the instructor's pedagogical practice; the corpus supports it across instructor reflection, public-facing writing, and live workshop delivery.

The label "hallucination-as-pedagogy" is my analytic coinage in this dissertation. The corpus contains the phenomenon — me reframing the tool's hallucinations as teachable moments across multiple registers from January 2024 through September 2025 — but the label itself emerges here, in the analytic work, not in the iterations' own materials. The same is true of the other two sub-claim names: "compression-as-curriculum-maturation" and "multi-channel teaching practice" are coined in this dissertation as analytic instruments rather than recovered from earlier sources.

D.1.3 Cross-validation as the rigor move

Each elaboration is anchored across multiple artifact types from multiple points in the timeline, with convergence treated honestly. The compression-as-maturation sub-claim anchors across the four iteration artifact sets; the multi-channel sub-claim anchors across eight documented contexts (whose heterogeneity is acknowledged in §D.3); the hallucination-as-pedagogy sub-claim anchors across four sources of which one (KT-THEMES-C5) is near-independent and three are Larissa's own surfacings across registers (§D.4).

The convergence across artifact types and across time is what distinguishes a defensible autoethnographic elaboration from an opportunistic observation. The dissertation does not claim more than the convergence supports.

D.1.4 What this chapter does not claim

I name three things this chapter does not claim, so that the scope of the autoethnographic elaborations is clear.

First, the elaborations do not replace the three principles. The DBR analysis of the main document is unaffected by what this chapter does; what this chapter does is sit alongside it and extend the analytic reach within each principle.

Second, the elaborations do not claim to be exhaustive. Analytic autoethnography of a multi-year multi-channel pioneer practice could surface other patterns within each principle that this dissertation does not develop. The three I elaborate are the strongest cross-validated patterns and the most useful for engineering education and HCI as scholarly fields.

Third, the elaborations do not claim that the patterns transfer automatically to other pioneer instructors of other emerging-technology curricula. Each pattern is a claim about my pioneer practice at the generative-AI site between 2023 and 2025. Whether the patterns travel to other sites is an empirical question for subsequent scholarship.

Sections D.2 through 4.4 develop each elaboration in turn.

D.2 Modularity · the architecture's stability seen from inside

The main document's first curriculum-design principle, **modularity**, names the structural property that lets independent course modules be reorganized, replaced, and recombined as the surrounding technology and audience change. This section elaborates the principle from the practitioner-pioneer's reflexive position. What I add is the autoethnographic record of **what stayed the same** across four iterations and eight contexts when the surface changed, and the analytic naming of the underlying pattern as **compression-as-curriculum-maturation**.

D.2.1 The architecture as the structural form of modularity

The four-theme architecture (Education, Industry, Ethics, Accessibility) is the structural form modularity takes in this curriculum. I built the themes into the course from Iteration 1 onward (DK-1.FG slides 5-9; FP-1's nine reflection questions distributed across the four themes; see also §A.4). Each theme operates as an independent module: any one theme could be taught without the others; the themes can be reordered; the tool-and-content list within each theme can be replaced as the field changes. Across the four iterations the architecture stayed intact while the surrounding curriculum shifted.

This is the autoethnographic observation that supplements the DBR principle. The DBR analysis says modularity emerged as a useful curriculum-design property; the autoethnographic supplement adds that the **specific architecture** that operationalized modularity (the four themes) is what carried the iterations through tool turnover and format compression.

D.2.2 What the corpus shows about architectural stability

The architecture held across four configurations:

Iteration	Format	Architecture in the corpus
Iteration 1 (CTD, Spring 2024)	15-week semester	ForeverGold deck (DK-1.FG) slides 5-9 name all four themes; FP-1 distributes nine reflection questions across them
Iteration 2 (GEEN 3830, Spring 2025)	15-week semester	Syllabus (SY-2) weekly structure: Weeks 1-4 Education, 5-11 Industry, 12-15 Accessibility, Ethics cross-cutting
Iteration 3 (CEAS, Aug 2025)	5-day workshop	Workshop deck (DK-3) one-topic-per-day mapping to the theme clusters

Iteration	Format	Architecture in the corpus
Iteration 4 (GenAI Works, Sept 2025)	5-day workshop	Workshop deck (DK-4) preserves the DK-3 template with date headers updated

Beyond the four iterations, the architecture appears in the HCI summer 2024 guest series (HC corpus: ten decks named by theme — HC-EDU, HC-INDUSTRY, HC-ACCESS, HC-AUDIO with Ethics, etc.) and in the public-facing channels (Keep Up Newsletter episodes organized by theme cluster; CU RMACC webinar organized around the four themes; see §D.3 for the multi-context elaboration).

D.2.3 The sub-claim · compression-as-curriculum-maturation

The autoethnographic analysis surfaces a pattern within modularity that I name compression-as-curriculum-maturation. The pattern is that the same modular architecture compressed approximately six-fold in delivery time (from 15 weeks to 5 days) while preserving its structural form. The compression is not content loss; it is content distillation that became possible because the architecture itself was modular and stable.

The diagnostic for the pattern is **architectural stability across compression in delivery time**. Where curriculum compression in technology fields is sometimes described as expansion (growing into multi-course sequences) or loss (shedding content under time pressure), my iterations exhibit a third pattern: the architecture persisted at higher pedagogical density while the surrounding scaffolding adapted to the new format.

D.2.4 The format confound I take seriously

Iterations 3 and 4 were not only compressed versions of the semester course. They were also a distinct delivery format: a free public online workshop hosted on Luma and partnered with GenAI Works, delivered through YouTube to an audience that registered for a five-day commitment. The compression from 15 weeks to 5 days is therefore the joint product of two forces: (a) what I had

learned to identify as essential through two full-length iterations, and (b) what the online-workshop format itself requires (online learners commit to short, intense windows; the platform supports 5-day rather than 15-week structures; the GenAI Works partnership was framed as a workshop from the outset).

A reader who attributes the compression entirely to instructor maturation would overstate what (a) accounts for; a reader who attributes it entirely to format constraints would understate what the two semester iterations made possible. My claim is the joint one: the format made compression **possible** in a way it would not have been without external constraints, and the two full-length iterations made the compression **successful** (the architecture held, the topics persisted, the satisfaction ratings were high) in a way it would not have been from a workshop-first starting point.

D.2.5 Why this pattern emerges in pioneer practice

The pattern emerges because pioneer practice has a learning curve that non-pioneer practice does not. A pioneer instructor in a fast-moving technology field cannot consult a literature on what is essential to teach, because the literature does not yet exist. The pioneer learns what is essential through full-length iterations that surface what students engage with, what tools persist beyond a season, and what conceptual frameworks travel across topics. Once the pioneer has accumulated that knowledge, compression to a short format becomes possible because the pioneer knows what to keep.

This is why my workshops succeeded the semester courses in the timeline rather than preceding them. Iteration 1 was the discovery phase; I did not know in advance which themes, tools, and assignments would land. Iteration 2 was the consolidation phase; I tested the framework against a new student population and refined the module placement, the tool list, and the guest-speaker roster. Iterations 3 and 4 were the distillation phase; with the framework stabilized, I compressed to the workshop format without losing what mattered.

A non-pioneer instructor adopting an established curriculum would not need to do this. The

accumulated literature would tell them what to keep and what to drop. The pioneer's compression is therefore not just a smaller version of the semester course; it is the outcome of the pioneer's two-semester learning process.

D.2.6 Theoretical contribution

The literature on iterative curriculum design (McKenney and Reeves 2018) treats iteration as a method for refining design conjectures through structured cycles. The pattern this elaboration identifies is consistent with that literature in the sense that iteration was the mechanism through which the curriculum matured. The elaboration adds two specifics that the iterative-design literature does not commonly name.

First, the maturation pattern is compression rather than expansion. The pioneer's iterative learning produced a shorter, denser delivery, not a longer or more elaborated one. This contrasts with how curriculum maturation is sometimes described in technology fields, where successful curricula are said to grow into multi-course sequences. Both patterns are possible; mine is the compression pattern.

Second, the compression preserves architecture rather than content alone. The themes were stable across all four iterations; the tools changed across the iterations; the duration compressed; the architecture survived. This suggests that what pioneer instructors learn through iteration is the architecture, not the content list. The content list is replaceable, and is replaced as new tools emerge. The architecture is the durable pedagogical contribution.

The elaboration contributes the claim that compression-as-curriculum-maturation is a documentable pattern in pioneer instructor practice at fast-moving technology fields, and that the diagnostic for the pattern is architectural stability across compression in delivery time. The claim sits as a sub-claim within the modularity principle the main document develops.

D.2.7 Guidelines for instructors entering fast-moving technology fields

The compression pattern documented in this elaboration is not a recipe; it is a posture that other pioneer instructors of emerging-technology curricula may find useful. Four guidelines drawn from the cross-iteration record:

- **Plan to run the curriculum at full length before compressing it.** The pioneer's first iteration is the discovery phase, and Iteration 2 is the consolidation phase (§C.3.8). Compression to a workshop format becomes possible only after architectural stability emerges through full-length delivery. Starting with a compressed format risks omitting what the longer iterations would have surfaced as essential.
- **Treat the conceptual architecture as durable; treat the tool list as disposable.** Themes stay. Tools turn over. The Midjourney Self-Portrait Assignment recurred across all four iterations under three generations of image-generation tools. The architecture absorbed the turnover; an architecture organized around specific tools would not have.
- **When compressing, drop the lecture-discussion-lab-assignment loop, not the topics.** The topics are what learners came for. The pacing structure is what semester length supports and workshop length does not. In the workshops, live demonstration takes the explanatory load that lecture-discussion cycles carried in the semesters.
- **Expect roughly a six-to-one compression ratio.** Iterations 3 and 4 documented this ratio from the fifteen-week semester to the five-day workshop. Pioneer instructors planning similar moves can use this as a rough planning anchor, recognizing that the ratio is empirical to my case rather than universal.

D.3 Learner Choice · dialogue with informants across channels

The main document's second curriculum-design principle, **learner choice**, names the property that lets learners select among tools, formats, and topics within the curriculum's structure. This section elaborates the principle from the practitioner-pioneer's reflexive position. What I add

is the autoethnographic record of **how learner choice operated across multiple contexts and audience types** — and the analytic naming of the underlying pattern as **multi-channel teaching practice**.

D.3.1 Learner choice as a multi-context phenomenon

The DBR analysis develops learner choice as a within-classroom property: students in a course choose among assignment formats, tool variants, presentation modes. The autoethnographic supplement extends the principle outward. My pioneering practice did not unfold in a single classroom. It unfolded across multiple delivery contexts simultaneously, and learner choice operated **between** contexts (which channel learners encountered the curriculum through) as well as **within** contexts (what they did inside each).

This wider reading of learner choice is what the autoethnographic position makes possible to see. The DBR within-classroom view of learner choice is correct as far as it goes; the autoethnographic view adds that the same four-theme architecture was simultaneously available through eight contexts of varying duration, audience, and depth, and that the heterogeneity of contexts is itself a form of learner choice.

D.3.2 The eight contexts documented

The architecture's appearance across these eight contexts is documented as follows. The trajectory of the four iterations across the six concurrent delivery channels is visualized as the opening figure of Appendix C (§C.1, Figure 1); the eight contexts of learner choice intersect that timeline, with the for-credit semester courses and the compressed online workshops appearing along the iteration spine, and the HCI graduate series, K-12 outreach, public-facing reflection, and federal-research webinar appearing in the other four channels.

Context 1 · Iteration 1 (CTD undergraduate course, Spring 2024)

The CTD pilot at the ATLAS Institute, fifteen weeks, approximately twenty-five students.

Context 2 · Iteration 2 (Mixed Engineering undergraduate course, Spring 2025)

GEEN 3830-001 Special Topics, fifteen weeks, approximately twenty-five students of mixed engineering backgrounds.

Context 3 · Iteration 3 (CEAS-sponsored online workshop, August 2025)

First cohort of GenAI in Five, five days at one hour per day, 129 live attendees out of 411 registered (LR-3).

Context 4 · Iteration 4 (GenAI Works partnership, September 2025)

Second cohort of GenAI in Five, YouTube delivery to a globally-distributed audience. Approximately 4,731 registered on the Luma event page (LR-4); the Day 1 broadcast drew 2,654 live online participants (LR-4.D1); Luma feedback aggregated to 256 ratings averaging 4.2 of 5 (LF-4).

Context 5 · HCI summer 2024 guest series

Ten guest lectures delivered for an HCI course at CU Boulder, covering the four themes across HCI applications.

Context 6 · K-12 outreach

STEAM Festival mural workshop with 50+ Colorado children using DALL-E 2; UW KidsTeam and Youth Advisory Board co-design sessions with elementary, middle-school, and high-school students.

Context 7 · LinkedIn newsletter and AI by Hand podcast

Three episodes of Keep Up Newsletter and two episodes of Keep Up Podcast.

Context 8 · Federal research webinar via CU RMACC and NAIRR Pilot

56-minute webinar on "Unleashing Creativity with Generative AI."

A note on what "eight contexts" should and should not be read as. The eight vary substantially in duration, audience size, depth of engagement, and institutional embedding. Two are

fifteen-week semester courses with ~45 hours of formal class time each; two are five-day workshops with ~5 hours; one is a ten-deck graduate guest-lecture series; one is a single-event K-12 outreach activity with continuing children's-research collaboration; one is a newsletter-plus-podcast pair; one is a single 56-minute federal-research webinar. Treating these as comparable units would overstate their similarity. What the elaboration documents is the appearance of the four-theme architecture across this range of contexts, not equivalent depth of engagement at each.

D.3.3 The sub-claim · multi-channel teaching practice

The autoethnographic analysis surfaces a pattern within learner choice that I name multi-channel teaching practice. The pattern is that the practitioner-pioneer's practice operates as a network of simultaneous delivery contexts rather than as discrete classroom delivery, and that learner choice in such a configuration extends to choice of context itself.

I claim that this practitioner-pioneer's practice at the technological frontier is most accurately characterized as networked multi-channel engagement rather than as discrete classroom delivery. The pioneer is not a single-classroom teacher who happens to engage occasionally with adjacent audiences; the pioneer is a multi-channel practitioner whose classroom work is one node in a wider network of pedagogical delivery.

I hedge the generalization carefully. Whether non-pioneer instructors of generative-AI courses also operate across multiple contexts, and whether the multi-channel pattern is specific to pioneer practice or is a more general feature of contemporary technology pedagogy, are empirical questions for subsequent comparative work. My evidence base is one practitioner over two and a half years; the claim is documented in this single case and is offered as a starting point for the comparative work that would test its generality.

D.3.4 Dialogue with informants beyond self · the evidence

Learner choice presupposes that learners' voices reach the curriculum and shape it. The autoethnographic supplement supports this through dialogue-with-informants evidence at multiple

scales:

- **Named undergraduates in semester iterations.** Twenty-three named student prompt-engineering outputs appear in DK-1.W01 alone, including Ashley Stafford and Ethan Cuenca. Six student teach-out presentations from Iteration 2 (SP-2 series). The Iteration 1 Final Project (FP-1) produced student-level four-theme syntheses.
- **Online learners in compressed iterations.** Twenty-nine Luma feedback responses (LF-3) carry evaluative comments; the Iteration 3 audience composition (LR-3) shows 65% students and 70% expressing interest in an AI Master's program.
- **K-12 children in the cross-iteration corpus.** The UW KidsTeam research (KT corpus) surfaced children's and teens' independent reasoning about generative AI in schools.
- **Journalists writing about the work.** Aspen Public Radio (AP-2024-05-16) quoted me and my Iteration 1 student Ashley Stafford one week after Iteration 1 ended.
- **Public-facing audiences.** The Keep Up Newsletter (KN-EP1 publicly credits a student in KN-EP3-Q1, "I learned about Soundful from one of my students during class") and federal-research audience via the CU RMACC webinar.

The convergence across types — named students, online learners, K-12 children, journalists, public audiences — is what makes dialogue-with-informants-beyond-self a substantively documented criterion across the four-theme architecture rather than an instructor's stray claim.

D.3.5 The Cuenca-to-Soundful flow as concrete instance

One named case anchors how learner choice flows back to shape the curriculum. Slide 13 of the Iteration 1 Week 1 deck (DK-1.W01) names Ethan Cuenca as one of the twenty-three students whose prompt-engineering outputs are displayed in the opening session. The Weekly Updates Prelim Document notes that one student presented on the GenAI music tool Soundful, and that I incorporated the student's discovery into the class assignment. Keep Up Newsletter Episode 3 (KN-EP3-Q1) publicly credits the flow: "I learned about Soundful from one of my students

during class." Soundful subsequently appeared in the Iteration 2 curriculum (CV-2), the Iteration 3 workshop deck (DK-3), and the Iteration 4 workshop delivery (TR-4.D3 carries Soundful as part of the Audio-generation day).

I want to be precise about what this is and is not: it is one well-documented case, not a pattern with multiple documented instances at the same level of precision. Other student tool discoveries shaped the curriculum across the iterations (the Iteration 1 Final Project, the Iteration 2 teach-out presentations, the Iteration 3 Luma feedback all surfaced tool suggestions), but they are not documented as precisely as the Cuenca-to-Soundful flow. The single named case is offered as a concrete instance of a broader phenomenon I observed but did not document case-by-case.

D.3.6 Theoretical contribution

The engineering-education and HCI literatures on instructor practice mostly take the single classroom as the unit of analysis. Multi-classroom comparative studies and longitudinal studies of single instructors across multiple courses are familiar; the multi-channel framing I propose is less common.

The contribution has practical implications for the engineering-education literature, calibrated to what this single-case study supports. It suggests that documenting pioneer instructor practice in a specific case requires documenting the full channel network rather than the classroom alone. It suggests that the stability of an instructional architecture, where multiple channels exist, is plausibly measured by cross-channel presence rather than by within-classroom outcome measures alone. It suggests that the K-12 outreach work, the public-facing writing, and the federal-research engagement that pioneer instructors often pursue alongside their classrooms are not extracurricular additions in cases like mine; they are part of the pioneering work and warrant documentation as such.

D.3.7 Multi-channel practice as engaged scholarship

Boyer's (1990) framework in **Scholarship Reconsidered** distinguishes four kinds of scholarship: discovery (traditional research), integration (synthesis across fields), application (engaged practice), and teaching (the pedagogical work itself). The multi-channel teaching practice documented in this elaboration maps cleanly onto Boyer's scholarship of engagement and the scholarship of teaching: the K-12 outreach (ST-MURAL, KT corpus), the public-facing newsletter and podcast (KN-EP series, KP-EP series), and the federal-research webinar (WB-2026-03-03) are not extracurricular to the dissertation's contribution; they are the scholarship of engagement in operation.

The Boyer mapping helps clarify what kind of contribution this elaboration is. It is not a contribution to the scholarship of discovery in the conventional sense (no controlled experiment, no novel measurement). It is a contribution to the scholarship of teaching and engagement: a documented record of how a pioneer instructor's pedagogical work propagates across multiple audiences and platforms, framed as scholarship rather than as an instructor's professional activity.

D.3.8 Guidelines for documenting and sustaining multi-channel teaching practice

The eight-context network documented in this elaboration is not unusual for pioneer instructors; what is unusual is documenting it as part of the work. Four guidelines drawn from the cross-channel record:

- **Plan the channel network alongside the classroom curriculum.** K-12 outreach, public-facing writing, podcasts, federal-research webinars, and grad-course guest lectures are not extracurricular additions to pioneer instructor practice; they are where the curriculum's architecture is tested against audiences beyond enrolled students. The 2024 calendar shows six channels operating simultaneously; that density is not exceptional for pioneer work.
- **Use cross-channel feedback as a validity check on the curriculum.** Children at UW KidsTeam (KT-THEMES) surfacing the same themes as enrolled undergraduates is

independent evidence that the architecture holds. Journalists framing the work in theme-consistent terms (AP-2024-05-16) is another. The cross-channel resonance is the stability test.

- **Make the channels feed each other.** Student tool discoveries in the classroom can be publicly credited in newsletters (KN-EP3-Q1 publicly credited an Iteration 1 student for Soundful). Guest speakers from one channel can re-appear in another with different topics (Nolan Brady in Iteration 1 and Iteration 2; Tom Yeh in Iteration 2 and Iteration 4). The relationships make the network coherent.
- **Document the channel network as part of the practice.** Scholarship that takes only the classroom as the unit of analysis will under-describe what pioneer instructors do. Documenting the channel network is one way to make pioneer practice legible to the engineering-education and HCI literatures.

D.4 Continuous Feedback · the reflexive loop in practice

The main document's third curriculum-design principle, **continuous feedback**, names the property by which surveys, reflections, and learner responses guide course revision across iterations. This section elaborates the principle from the practitioner-pioneer's reflexive position. What I add is the autoethnographic record of **how feedback operated at multiple timescales** — within a single lesson, across an iteration, and across iterations — and the analytic naming of the most concrete in-the-moment instance as **hallucination-as-pedagogy**.

D.4.1 Continuous feedback as a multi-timescale phenomenon

The DBR analysis develops continuous feedback as a between-iteration property: surveys and reflections at the end of each iteration inform the next. The autoethnographic supplement adds two further timescales the DBR analysis does not naturally surface:

- **In-the-moment** · the instructor's feedback loop with the tools and the learners during a

single lesson. When a tool produces an unexpected output (a hallucination, an accessibility limitation, a surprising affordance), the instructor's immediate reframing of that output into a teachable moment is the moment-to-moment form of continuous feedback.

- **Within-iteration** · feedback as it accumulates across the weeks of a semester course or the days of a workshop. The Weekly Updates Prelim Document (WU-1 series) is the densest documentation of this timescale in the corpus.
- **Across-iteration** · feedback as it shapes successive iterations of the course. The DBR-analyzed timescale.

The DBR analysis of the main document covers the third timescale (across-iteration). The autoethnographic supplement extends to the first two.

D.4.2 What the corpus shows about continuous feedback across timescales

The principle operates across all three timescales in the corpus:

Timescale	Evidence
In-the-moment	WU-1.W01-Q1 captures my first-week observation of ChatGPT-quiz hallucinations and the immediate reframing as teachable. TR-4.D1-Q3 captures the live workshop passage where I name a hallucination, show the concrete example (no body in the shoes), normalize the phenomenon, and convert the encounter into a re-prompting practice.
Within-iteration	The fifteen-week Weekly Updates Prelim Document (WU-1.W01 through WU-1.W15) is the most complete documentation of weekly feedback shaping next-week instruction in any iteration. The "Learned" sections are first-person reflective data captured at the time of teaching.

Timescale	Evidence
Across-iteration	The Iteration 1-to-Iteration 2 tool turnover (DeepSeek added in Iter 2 Week 3 within weeks of public release; AI Agents as a new full module; structural reshuffles documented in §C.6) is the principal cross-iteration feedback evidence. The Iteration 3-to-Iteration 4 compression also operates at this timescale (workshop deck DK-4 preserved the DK-3 template with date headers updated).

D.4.3 The sub-claim · hallucination-as-pedagogy

The autoethnographic analysis surfaces a pattern within continuous feedback at the in-the-moment timescale that I name hallucination-as-pedagogy. The pattern is that generative-AI hallucination, dominantly framed in the technical literature as a system limitation, is also a productive pedagogical phenomenon: the instructor's moment-to-moment encounter with hallucination becomes a teachable opportunity to expose the limits of the technology, invite human verification as a counter-practice, and open space for ethics and accessibility conversations.

The label "hallucination-as-pedagogy" is my analytic coinage in this dissertation. The corpus contains the phenomenon — me reframing the tool's hallucinations as teachable moments across multiple registers from January 2024 through September 2025 — but the label itself emerges here, in the analytic work, not in the iterations' own materials.

D.4.4 Source 1 · Instructor reflection in Iteration 1 Week 1

My first encounter with hallucination as a teachable moment is recorded in the Weekly Updates Prelim Document for Iteration 1 Week 1 (WU-1.W01). I wrote in the "Learned" entry for that week (WU-1.W01-Q1, verbatim):

> "I learned that some of the multiple choice quizzes generated by ChatGPT were not correct and had hallucinations."

The note is brief and dated to the very first week of the very first iteration. It captures what I noticed at the time: that the tool I had given students to generate quiz material had produced factual errors, and that the errors were not a deal-breaker for the pedagogy but were themselves something to teach about. The persistence of hallucination as a topic across subsequent iterations (sources 2 through 4 below) indicates that the teachable-moment reframing took hold from this first observation forward.

D.4.5 Source 2 · K-12 children's observation at UW KidsTeam (co-design context)

Six months after Iteration 1 closed, in July 2024, I joined the University of Washington KidsTeam group and Youth Advisory Board on co-design sessions about generative AI in schools. The children and teens were not coached toward any framework I had developed, but I want to characterize the corroboration carefully: I was a participant in the co-design sessions, not an external observer. The children's themes were elicited by the KidsTeam methodology (which the UW team led, not me) and were articulated in their own terms, but the co-design context was collaborative, not independent of my presence in the room.

What the children's articulation provides, then, is corroboration that is **near-independent** of my own curriculum: the methodology was UW's, the children were not my students, and the analysis was the UW team's. What it is not is fully independent observation. With that caveat noted, the KidsTeam research themes document (KT-THEMES) records the children's challenges with generative AI in their own catalogued form. The hallucination challenge appears as the fifth named challenge (KT-THEMES-C5, verbatim):

> "Hallucinations (such as images produced with a third arm)"

The same document records a broader pattern of children's challenges with the tools alongside the hallucination observation: "Generative AI being banned from schools," "Cheating," "AI is data fed into a computer and doesn't yet know everything," and "Lack of emotion and human connection" (KT-THEMES-C1 through KT-THEMES-C4 respectively). The hallucination obser-

vation is therefore not an isolated complaint; it sits within a structured list of named challenges the children produced through the UW KidsTeam co-design protocol.

The convergence with my own observation is the strongest single piece of near-independent evidence in the corpus.

D.4.6 Source 3 · Public-facing reframing in Keep Up Newsletter Episode 1

In April 2025, in the first episode of my Keep Up Newsletter on LinkedIn (KN-EP1), I addressed hallucination publicly as part of the "Lessons from Training" section of the episode's running-and-training metaphor (KN-EP1-Q1, verbatim):

> "Expect variable results, occasional hallucinations; persistence improves prompting skills; join community groups"

The phrase makes the pedagogical reframing explicit, and the surrounding section makes the reframing structural. Hallucinations are not a fault that should make the reader abandon the tools; they are expected behavior alongside variable results, and the appropriate response is persistence in prompting and engagement with community knowledge. The newsletter audience is a general LinkedIn readership, not a classroom; the reframing is therefore not classroom-specific but a posture I have generalized from the classroom and offered as public guidance.

D.4.7 Source 4 · Live workshop delivery in Iteration 4 Day 1

In September 2025, in the first day of the Iteration 4 GenAI Works workshop, hallucination appears in the live teaching delivery (TR-4.D1-Q3, verbatim):

> "And you can also see these different hallucinations that are going on. So there aren't even bodies in these shoes. And so it's really interesting how these different tools are used. Sometimes I get some really really great image generation and it just depends on what you want to use it for, but other times there are a lot of hallucinations that you can get within these specific tools. And so once again, you just have to make sure and take a look at everything. And even if you do get a hallucination on the first time, it doesn't mean give up. You can keep inputting and inputting and inputting the prompts into into the generator and see if it changes.

> And so I would I would continue to keep working with the specific tool and if you're still not getting a result that you like, then try a different tool."

The passage stages four moves in roughly forty seconds of teaching: **naming** (the technical term "hallucination"), **showing** (the concrete shoes-without-bodies example), **normalizing** (there are "a lot of hallucinations" learners will encounter), and **iterating** (re-prompt, keep prompting, switch tools). The four moves together are the operational form of hallucination-as-pedagogy. The passage is also the most recent of the four sources for this elaboration; by Iteration 4, the posture was a stabilized component of my workshop curriculum.

D.4.8 The convergence across sources

The four sources span twenty months (January 2024 through September 2025) and four artifact types. Three of the four sources (WU-1.W01-Q1, KN-EP1-Q1, TR-4.D1-Q3) are surfacings of the phenomenon by me, at different points in my practice and across different registers: contemporaneous reflective journaling, public-facing instructional writing, and live workshop delivery. They document the **stability** of my instructor reframing over time, across registers, but they are not independent observations. The fourth source (KT-THEMES-C5) is the children's own observation in a UW KidsTeam co-design context, and is the corroboration nearest to independence.

What the convergence supports is two claims at different strengths. The stronger claim is that **my own framing of hallucination as a teachable phenomenon stabilized across my practice's registers**, from private reflection to public writing to live workshop. The weaker claim, supported by one near-independent observation, is that **children in adjacent contexts surface the same phenomenon as a concern**. Both claims are part of the elaboration.

D.4.9 Theoretical contribution

The technical literature on generative AI predominantly treats hallucination as a system property to be measured and reduced. Bender, Gebru, McMillan-Major, and Shmitchell (2021) provide the canonical formulation in "On the dangers of stochastic parrots," which frames hallu-

cination among the broader harms of large language models that the field should design against. The technical lineage is appropriate to its goals: building better systems.

The pedagogical literature on how to teach with generative-AI tools has not yet developed a corresponding framework for hallucination as a teaching resource. The empirical question this elaboration addresses is not whether hallucination is a system property (it is) or whether system properties should be reduced (they should). The question is what an instructor does with hallucination when it appears in front of learners.

The elaboration contributes the claim that hallucination, encountered in instructional settings, is most productively framed as a teachable phenomenon that exposes the limits of the technology, invites human verification as a counter-practice, and opens space for ethics and accessibility conversations. The contribution is not a curriculum module or a teaching technique; it is a posture an instructor can adopt and that learners across age groups appear to find legible. The claim sits as a sub-claim within the continuous-feedback principle the main document develops.

D.4.10 Guidelines for educators teaching with generative-AI tools

The four moves I stage in TR-4.D1-Q3 (name, show, normalize, iterate) can be transposed by other instructors. Five guidelines, drawn from the cross-source convergence:

- **Do not avoid the live encounter with hallucination.** Working only with carefully-curated tool outputs hides the most teachable feature of the technology. Let learners see hallucinations as they appear during a workshop or class.
- **Name the phenomenon by its technical term.** Children at UW KidsTeam used the word "hallucinations" (KT-THEMES-C5); undergraduates in Iteration 1 encountered the term in Week 1 reflections (WU-1.W01-Q1); online learners met the term in public-facing writing (KN-EP1-Q1). The vocabulary is legible across age groups. Softening or apologizing for the term weakens the lesson.
- **Convert the encounter into a re-prompting practice.** The classroom move I demon-

strate in TR-4.D1-Q3 is concrete: keep prompting; vary the prompt; if the tool will not produce what you want, switch tools. This converts a moment of apparent system failure into a piece of learner skill.

- **Use hallucination as the bridge to ethics and accessibility themes.** What hallucinations get wrong about human bodies (the "third arm" example) opens accessibility conversations about whose bodies these tools can represent. What hallucinations get wrong about facts opens ethics conversations about verification and trust.
- **Treat public-facing AI-literacy work the same way.** Newsletters, podcasts, and federal-research webinars can productively name hallucination as expected behavior. The audience is users who will encounter it, not consumers who should be protected from it.

D.5 Contributions

This section names the contributions the dissertation makes to scholarship, organized around the three curriculum-design principles and their autoethnographic elaborations. The contributions are of three kinds: methodological, substantive, and pedagogical. I take each in turn, name its claim, locate its evidentiary base, and indicate the literatures to which it speaks. I close with limitations and future directions.

D.5.1 Methodological contribution · supplementing design-based research with analytic autoethnography

My methodological contribution is the demonstration that analytic autoethnography (Anderson 2006) can supplement design-based research at a generative-AI pedagogy site, deepening each of the three curriculum-design principles the main document develops by surfacing the practitioner-pioneer's reflexive account of how the principles emerged from her practice and what they reveal at a closer analytic level.

The methodological contribution is the supplementation move itself, not either methodology in isolation. DBR is a mature methodology; analytic autoethnography is a twenty-year-old

framework; neither is novel. What is new in the engineering-education and HCI literatures is the combined deployment of the two at a fast-moving emerging-technology pedagogy site, with the autoethnographic supplement reading the same artifact corpus that the DBR primary analysis reads but surfacing a different layer of analytic detail within the same principles.

Mapped to Wobbrock and Kientz's (2016) taxonomy of research contributions in HCI, the methodological contribution is what they call an **opportunistic contribution** (deploying an established methodological supplementation in a new configuration), folded together with a **survey contribution** (the per-iteration record itself constitutes a survey of how a pioneer instructor's practice unfolds at the GenAI moment). It is not an **empirical** contribution in the controlled-experiment sense; the dissertation does not claim that pattern.

D.5.2 Substantive contribution · the three principles enriched by autoethnographic elaboration

My substantive contribution is the three curriculum-design principles **modularity**, **learner choice**, and **continuous feedback** as the main document develops them, **enriched by three autoethnographic elaborations** that this appendix surfaces. The principles are the practitioner-facing outputs the field can adopt; the elaborations are the scholar-facing outputs that explain how the principles emerged and what they reveal at a closer analytic level.

Principle	Autoethnographic elaboration (sub-claim)
Modularity	The four-theme architecture's documented stability under tool turnover and format compression. Compression-as-curriculum-maturation as the analytic name for the pattern (§D.2).
Learner Choice	Dialogue with informants across eight contexts of varying duration, audience, and depth. Multi-channel teaching practice as the analytic name for the pattern (§D.3).

Principle	Autoethnographic elaboration (sub-claim)
Continuous Feedback	The reflexive feedback loop the instructor maintained with tools and learners. Hallucination-as-pedagogy as the analytic name for the most concrete in-the-moment instance (§D.4).

The substantive contribution is neither the principles alone (DBR could surface them) nor the elaborations alone (autoethnography could surface them). It is the two layers working together: principles that adopters can use and elaborations that explain to scholars what the principles' adoption requires of the practitioner.

D.5.3 Pedagogical contribution · the four-theme architecture's documented stability

My pedagogical contribution is the empirical record of how the four-theme curriculum architecture (Education, Industry, Ethics, Accessibility) functions as a stable organizing structure under tool turnover, audience shift, and delivery-format compression. The contribution is the **stability record**, not the **category names** (§A.4.4 acknowledges that adjacent AI-literacy frameworks use overlapping category names; the four-theme list itself is not novel).

The stability is documented along three dimensions:

- **Across iterations and tool turnover.** The architecture held across four iterations spanning two and a half years and at least three generations of underlying tools (image generation from DALL-E to Midjourney to Nano Banana; large language models from ChatGPT to DeepSeek to Claude 3.7; new capabilities including AI Agents and accessibility-targeted tools like Be My AI).
- **Across delivery contexts.** The architecture held across eight contexts (§D.3) from undergraduate semester courses to online workshops to K-12 outreach to federal-research webinars.
- **Across audience age groups.** The architecture held against learners from elementary-school age through graduate students and professional federal-research audiences, with

cross-channel resonance documented in §A.5.

The architecture is testable. A subsequent instructor of an emerging-technology curriculum could adopt this architecture, deliver under it, and report whether it holds against their cohort, their tools, and their institutional context. The architecture is small enough to remember (four themes), explicit enough to operationalize (each theme has a named scope; ethics is positioned cross-cuttingly), and flexible enough to accommodate substantial tool turnover (my own corpus documents this).

The pedagogical contribution speaks to the engineering-education and HCI pedagogy literatures and to the broader AI-literacy literature emerging in 2024 through 2026. It is not the only valid framework for generative-AI pedagogy. It is a framework whose architectural stability has been empirically documented across the configurations named above.

D.5.4 Consolidated guidelines for educators teaching generative AI

Anderson's (2006) framework for analytic autoethnography licenses pedagogical-implications work as a legitimate genre of contribution: the insider's analytic position is supposed to produce insight that travels to adjacent practitioners. The guidelines below are that kind of contribution. They are not the autoethnographic elaborations themselves (§D.2 through §D.4); they are the educator-facing distillation of what those elaborations suggest for practice, mapped to the three principles. They are offered with the appropriate hedging: the single-case scope of this dissertation (§D.5.5) means the guidelines are starting points for adopters in adjacent settings rather than empirical prescriptions.

Across the three principles, five top-line guidelines emerge for educators planning, delivering, or sustaining a generative-AI curriculum:

- **Build the curriculum around a small set of stable themes that can absorb tool turnover.** This is the modularity guideline: small enough to remember, explicit enough to operationalize, flexible enough to accommodate the next tool the field releases. The four

themes that worked here (Education, Industry, Ethics, Accessibility) are one starting set; other configurations are plausible.

- **Iterate at full length before compressing.** This is the modularity-and-maturation guideline: the five-day workshop format is the outcome of two semester-length passes, not the starting point. Pioneer instructors compressing to a workshop without prior full-length delivery will likely omit what the longer iterations would have surfaced as essential.
- **Run the curriculum across multiple contexts deliberately.** This is the learner-choice guideline: K-12 outreach, public-facing writing, podcasts, and federal-research webinars are not extracurricular. They are where the architecture is tested against audiences beyond enrolled students, and where cross-channel resonance becomes evidence that the architecture holds.
- **Credit students publicly when their tool discoveries shape the curriculum, and document the flow.** This is also a learner-choice guideline: the Ethan Cuenca to Soundful case (§D.3.5) is one named instance. The student-to-instructor tool flow is part of pioneer practice and warrants documentation as such.
- **Treat hallucination, error, and surprise as the teachable centerpieces, not the failures to hide.** This is the continuous-feedback guideline: what the technical literature frames as system limitations are what an instructor in front of learners can convert into the most legible moments of the curriculum. Name the phenomenon, show it, normalize it, convert it into a re-prompting practice.

These guidelines are calibrated to pioneer instructor practice in fast-moving technology fields. Whether they travel to other domains (mature curricula, slower-moving fields, multi-instructor programs) is an empirical question for subsequent scholarship, addressed in §D.5.7 Future directions.

D.5.5 Limitations

I name three limitations of the contributions claimed above.

Single-instructor, single-program scope. My dissertation documents one instructor's pioneering practice at one research university. The principles, the autoethnographic elaborations, and the guidelines are claims about pioneer practice at the generative-AI site between 2023 and 2025. They do not generalize automatically to other pioneers, other sites, or later periods when the field has matured beyond pioneer-entrant status. Whether they travel is an empirical question for subsequent scholarship.

Asymmetric reflective data across iterations. My contemporaneous structured reflective journaling is concentrated in Iteration 1 (WU-1.W01 through WU-1.W15). Iterations 2, 3, and 4 draw on my public-facing reflective writing (KN-EP series across the 2025 iterations, KP-EP series in May 2025, WB-2026-03-03 in March 2026) as the supplementary reflective base, within an acknowledged retrospective-public frame that analytic autoethnography permits (§B.6.3). The reflective base for the later iterations is therefore narrower than for Iteration 1 but not absent.

Researcher's investment in the work. I have a stake in the work succeeding and in the autoethnographic elaborations landing as contributions. Analytic autoethnography does not require neutrality, and §B.7.4 names my investment explicitly. The trustworthiness moves I make in response (multi-source anchoring of each elaboration, named limits of the data, dialogue with chair and committee) are described in §B.6 and §B.7. They do not eliminate the investment; they discipline it.

D.5.6 Why this matters now

The dissertation's timing is part of its contribution. The four iterations span the post-ChatGPT period when generative-AI pedagogy emerged as a field, and the dissertation is, in its specific autoethnographic configuration, one of the empirical records of what the emergence looked like from inside an instructor's practice. Three reasons this matters now, beyond the substantive contributions named in §D.5.2.

First, the empirical record of **what an instructor learned in 2024 and 2025** will be retrospectively important to AI-pedagogy scholarship as the field matures. Reconstructions written

years later will not have access to contemporaneous artifacts of the kind preserved here: the Weekly Updates Prelim Document captured in real time as the tools were released; the workshop transcripts capturing teaching delivery while the audience-engagement patterns were still novel; the public-facing newsletter and podcast written as the instructor was still figuring out what to teach. The historical record value of this corpus increases with time.

Second, the methodological framing offers an alternative to the controlled-experiment paradigm that dominates engineering-education research on emerging-technology pedagogy. Pioneer instructors of fast-moving fields cannot, as a practical matter, run the controlled studies the dominant paradigm asks for: the technology changes faster than the IRB approves, the comparison condition does not exist (there is no established curriculum to compare against), and the timeline of the research mismatches the timeline of practical relevance. Analytic autoethnography of the kind this dissertation performs is one way the field can do legible scholarship at the timescale the technology demands.

Third, the educator-facing guidelines (§D.5.4) are time-sensitive contributions. Educators teaching generative AI in 2025 and 2026 face design choices the field has not yet stabilized answers to: how to handle hallucination in real time, whether to compress to workshop format or hold semester length, whether to extend across multiple contexts. The guidelines drawn from this dissertation are oriented to those choices and travel best while the choices are still open.

D.5.7 Future directions

I name three directions for subsequent scholarship that the dissertation opens.

Comparative autoethnographies of pioneering instructors. Subsequent analytic autoethnographies of other early-entrant generative-AI instructors at other institutions would test whether the three autoethnographic elaborations I claim travel beyond my case. A small comparative collection of pioneer-instructor analytic autoethnographies would also begin to build the empirical record that the field currently lacks.

Longitudinal tracking of students into practice. My corpus documents student-to-

instructor flow at a specific scale (Ethan Cuenca and Soundful, the named case in §D.3.5) and student-public engagement at a specific scale (Ashley Stafford quoted in Aspen Public Radio, AP-2024-05-16). Longer-term tracking of students from pioneer GenAI courses into industry, education, or research roles would extend the documentation of how pioneering instruction propagates.

Curriculum architecture portability studies. The four-theme architecture is testable against subsequent instructor cohorts. A study in which two or three subsequent instructors of generative-AI courses adopt the architecture, deliver under it, and report on its hold against their cohorts would empirically test portability (the next claim past the stability record this dissertation documents in §D.5.3).

These directions are not promised by my dissertation. They are openings the dissertation creates for scholarship that follows it.

D.6 Artifact inventory

In this section I name the artifact corpus that the dissertation rests on and the conventions by which the chapters cite into it. The full master index is the evidence table presented as Appendix E; in this section I summarize and orient the reader to that document.

D.6.1 Volume summary

The artifact corpus comprises roughly **105,000 words** of iteration-specific text data across the four iterations and approximately **35,000 words** of cross-iteration material, for a total of approximately **140,000 words** of contemporaneous and instructor-present text data. Per-iteration headline volumes: ~25,000 words for Iteration 1 (including the fifteen-week Weekly Updates Prelim Document); ~20,000 words for Iteration 2; ~5,000 words for Iteration 3; ~55,000 words for Iteration 4 (the five-day YouTube transcripts are the largest single source in the corpus).

The full per-artifact catalog, with IDs, periods, descriptions, and quote sub-IDs, appears as Appendix E. Every in-text citation in this dissertation maps to an entry in Appendix E, and that document is the source of truth for artifact provenance.

D.6.2 Metadata conventions

Each artifact in the corpus carries metadata that the evidence-table master index records.

- **ID**: hierarchical and self-describing, with an optional Q-suffix for quote-level citation.
- **Date**: exact where possible (e.g., AP-2024-05-16 for the Aspen Public Radio article dated May 16, 2024) and approximate where the artifact's date is inferred (e.g., RE dated c. Spring 2024 from internal "during this spring semester" language).
- **Iteration**: marked when the artifact is iteration-specific (1, 2, 3, or 4) or as "cross" when the artifact spans the iterations.
- **Description**: a one-line summary suitable for the evidence table's index.

Verbatim quotes are preserved with original punctuation. Where minor clarification is necessary for reader comprehension, `[sic]` marks are used. Where an artifact's text is paraphrased rather than verbatim (as is currently the case for the Aspen Public Radio article, AP-2024-05-16), the master index marks the paraphrase status explicitly and queues the verbatim extraction.

D.6.3 Limits of the corpus and what they imply for the analysis

I name two limits of the corpus that the analysis works within rather than against.

Asymmetric contemporaneous reflective data across iterations. Iteration 1 has the Weekly Updates Prelim Document (WU-1.W01 through WU-1.W15) as a fifteen-week structured reflective journal. Iterations 2, 3, and 4 do not have an equivalent contemporaneous structure. The cross-iteration reflective channels (KN-EP series, KP-EP series, WB-2026-03-03) carry my reflective writing across the later iterations within an acknowledged retrospective-public frame, and analytic autoethnography permits this configuration when the frame is named (§B.6.3). The reflective base for Iterations 2 through 4 is therefore narrower than for Iteration 1 but not absent.

Aspen Public Radio cited at source-level rather than sentence-level. AP-2024-05-16 is held as a source-level reference in the evidence-table. The article carries named quotes of

me and of my Iteration 1 student Ashley Stafford; the dissertation cites the article at source-level rather than at sentence-level.

The two limits are common in analytic autoethnographic work that surveys a multi-year practitioner-pioneer practice across multiple delivery channels. Naming them is what trustworthy autoethnography requires. Neither weakens the substantive contribution claims of the dissertation; they bound the evidentiary base on which those claims rest.

D.6.4 Cross-reference to Appendix E

The full per-artifact catalog, including all sub-IDs for Weekly Updates (WU-1.W01 through WU-1.W15), Luma responses (LF-3.R01 through LF-3.R29), and the quote-level Q-IDs the chapters cite, appears as Appendix E.

Appendix E

Evidence Table

Evidence Table

Master catalog of all artifacts cited in the supplementary chapters. Each artifact has a stable ID. Chapters cite by ID using inline parentheses, e.g., `(WU-1.W01)` or `(WU-1.W01-Q1)` for a specific verbatim quote.

The catalog is organized by iteration, then by artifact category within each iteration. Each iteration section opens with a brief overview of the iteration's institutional context and what its artifact corpus contains. Each category subsection opens with a summary of what kind of evidence the category provides for the chapters' analytic claims.

E.1 ID scheme

Format: `{category}-{iteration}.{location}` with optional `-Q{n}` for quote-level citation.

Prefix	Category
WU	Weekly Updates Prelim Document
CV	Canvas LMS export
SY	Syllabus
DK	Slide deck (lecture or workshop)
TR	Transcript (YouTube, podcast, webinar)

Prefix	Category
FP	Final Project artifact
SP	Student teach-out presentation
SO	Named student output
LF	Luma Feedback (one response)
LR	Luma Roster (aggregate)
AS	Administrative spreadsheet
RE	Research Impact Essay
KN	Keep Up Newsletter episode
KP	Keep Up Podcast episode
HC	HCI summer 2024 deck
KT	UW KidsTeam research item
ST	STEAM Festival artifact
AP	Aspen Public Radio article
AI	AI-IRT Seed Grant proposal
PR	Preliminary Exam document
WB	Webinar (CU RMACC)
STC	Storytelling Cartoonimator worksheet

E.2 Iteration 1 · CTD Pilot · Spring 2024

The first iteration ran from January 17 through May 1, 2024, as the CTD pilot at the ATLAS Institute with approximately twenty-five undergraduate students. This is the discovery phase of my pioneer practice: the curriculum was being built from scratch, the four-theme architecture was being stabilized through delivery, and my reflective journaling was the densest contemporaneous record of any iteration. The Iteration 1 corpus is the most analytically valuable single iteration's

data in the dissertation because it carries both the curricular structure as students experienced it and my week-by-week analytic reflection on what the teaching was teaching me.

E.2.1 Curricular artifacts

The curricular corpus for Iteration 1 documents what was taught, in what sequence, with what materials. It includes the full Canvas LMS export, the course-framing deck that established the four-theme architecture for the iteration, the Week 1 opening lecture with twenty-three named student outputs, and the Final Project assignment that operationalized the four themes through a Media Studies-style company-creation task.

CV-1 · Canvas LMS export, Spring 2024. 225 files; 15-week modules including 4 video recordings of guest lectures.

DK-1.FG · ForeverGold course deck, Spring 2024. 36 slides; the four-theme architecture (Education, Industry, Ethics, Accessibility) appears on slides 5 through 9.

DK-1.W01 · Week 1 opening lecture, January 17, 2024. 35 slides; 23 named student prompt-engineering outputs, including Ashley Stafford and Ethan Cuenca whose work persists across subsequent iterations.

FP-1 · Final Project Requirements, Spring 2024. Multi-tool Media Studies-style company-creation assignment with 9 reflection questions distributed across the four themes, including a dedicated ethics question on creative-job displacement.

E.2.2 Reflective artifacts

The reflective corpus for Iteration 1 is uniquely well-developed in the dissertation. The Weekly Updates Prelim Document is the only contemporaneous structured reflective journal in the corpus, kept across fifteen weeks with a consistent format ("Curriculum Development / Special Guests / Resources / Learned / Students Produced") per week. The "Learned" sections function as analytic reflexivity in Anderson's strict sense: each week records what the teaching had taught me, captured at the time the teaching happened. This is the densest contemporaneous reflective record any iteration produced.

WU-1.W01 · Weekly Update Week 1 (lecture date January 17, 2024). My first structured "Learned" reflection.

- **WU-1.W01-Q1**: "I learned that some of the multiple choice quizzes generated by ChatGPT were not correct and had hallucinations."
- **WU-1.W01-Q2**: "I learned that the students enjoyed creating the alter-ego images to describe themselves to the class."

WU-1.W02 through WU-1.W15 · Selected "Learned" entries from Weeks 2 through 15. Cited in Chapter C §C.2.8 as the arc-anchoring quotes across the semester. Full text in the Weekly Updates Prelim Document.

- **WU-1.W02-Q1**: Week 2 entry on Midjourney access and image-generation bias
- **WU-1.W03-Q1**: Week 3 entry on AI content detectors and institutional GenAI policy
- **WU-1.W05-Q1**: Week 5 entry on cross-industry GenAI application
- **WU-1.W06-Q1**: Week 6 entry on SORA and GenAI in finance
- **WU-1.W07-Q1**: Week 7 entry on deepfake ethics and music generation
- **WU-1.W14-Q1**: Week 14/15 entry on assistive technologies and Final Project execution

E.3 Iteration 2 · Mixed Engineering · Spring 2025

The second iteration ran from January 13 through April 30, 2025, as GEEN 3830-001 Special Topics in the College of Engineering and Applied Science with approximately twenty-five undergraduates of mixed engineering backgrounds. This is the consolidation phase of my pioneer practice: the four-theme architecture had stabilized through Iteration 1, the tool landscape had shifted substantially (DeepSeek, AI Agents, NotebookLM, Claude 3.7 all entered the curriculum), and the iteration introduced the student teach-out as a new pedagogical move. The Iteration 2 corpus is the strongest single source for the tool-turnover and module-placement-evolution patterns developed in Chapter C §C.6.

E.3.1 Curricular artifacts

The curricular corpus for Iteration 2 documents the consolidated mid-iteration form of the curriculum: a syllabus that names the four-theme architecture explicitly with weekly structure, the Week 1 prompt-engineering opener that preserved Iteration 1's structural commitment, and five iteration-specific lecture decks covering the new tools and themes that emerged between Spring 2024 and Spring 2025.

CV-2 · Canvas LMS export, Spring 2025. 240 files; 15 weeks of modules; 103 HTML pages; 6 video recordings of guest lectures.

SY-2 · Syllabus for GEEN 3830-001 Special Topics. Three-credit course; classroom DLC170; weekly structure naming the four themes; tool list including DeepSeek, Claude 3.7, Grok 3, Sora, HeyGen, Be My AI, ElevenLabs.

DK-2.JAN13 · Week 1 Prompt Engineering opener (January 13, 2025).

DK-2.JAN15 · Midjourney Self-Portrait Assignment deck (January 15, 2025).

DK-2.FEB05 · Future Wheel technique deck (February 5, 2025).

DK-2.MAR05 · Industry theme deck (March 5, 2025).

DK-2.MAR10 · Education theme deck (March 10, 2025).

E.3.2 Learner-facing artifacts

The learner-facing corpus for Iteration 2 includes student teach-out presentations, a new pedagogical move I introduced in this iteration. The teach-outs gave students an opportunity to research a topic of their choice and present it to the class, generating six recorded presentations preserved in the archive. The teach-out topics are distributed across the four themes and provide concrete evidence for the dialogue-with-informants-beyond-self criterion (§B.3.4).

SP-2.DEEPFAKE · Student teach-out: DeepfakeAI (Daniel Debretsion). Ethics topic with Industry implications.

SP-2.SINTRA · Student teach-out: Sintra.ai (Alt Style). Industry topic on AI-driven business workflow tools.

SP-2.HALEY · Student teach-out: Haley Phillips (two files).

SP-2.DAKOTA · Student teach-out: Dakota A.

SP-2.ROBOTICS · Student teach-out: The Integration of AI and Robotics. Industry topic on multi-modal AI in robotics.

E.3.3 Administrative artifacts

The administrative corpus for Iteration 2 documents the iteration's operational organization: the guest-lecturer rotation, the teach-out scheduling, and the final-presentation order. These spreadsheets are the operational trace of the curation work that Chapter C §C.6.3 develops as a pedagogical contribution in its own right; they show how the guest network and the student presentations were assembled into a coherent semester structure.

AS-2.GUESTS · Guest lecturer list spreadsheet. Records Tom Yeh (DeepSeek), Nolan Brady (NeuroImaging), Bobby Hodgkinson (NotebookLM), Justin Shacklette (GenAI for Software) as the iteration's central guests.

AS-2.TEACHOUT-DATES · Teach-out dates spreadsheet.

AS-2.TEACHOUT-INDEX · Teach-out slides index.

AS-2.FINAL-SCHED · Final presentations schedule.

E.4 Iteration 3 · GenAI in Five Cohort 1 · Aug 2025

The third iteration was the first compression of the curriculum from a fifteen-week semester to a five-day online workshop, running August 18 through August 22, 2025, sponsored by the College of Engineering and Applied Science (CEAS) and hosted on Luma. This is the first phase of distillation in my pioneer practice. The Iteration 3 corpus is the strongest single source for the workshop-format viability claim and for the learner-facing feedback that supports the compression-as-curriculum-maturation finding (§D.3).

E.4.1 Curricular artifacts

The curricular corpus for Iteration 3 is the workshop deck that compressed the full semester's content into a 22-slide, five-day structure with one topic per day (Monday image, Tuesday video, Wednesday audio, Thursday research, Friday human-centered AI).

DK-3 · Workshop deck, 22 slides; one topic per day across Monday through Friday.

E.4.2 Learner-facing artifacts

The learner-facing corpus for Iteration 3 is the principal triangulation data for the iteration. The Luma feedback survey produced twenty-nine responses with nine carrying text comments; the Luma participant roster documents the audience composition with unusual specificity (411 registered, 129 attended live, 65% students, 70% Master's-interested). Together, these data sources support the iteration's pedagogical-viability claim and surface the depth-versus-breadth tension developed in Chapter C §C.4.6.

LF-3 · Luma feedback aggregate. 29 responses; 22 fives, 6 fours, 1 two; average 4.69 of 5; 9 responses with text comments.

- **LF-3-Q1** (5★): "Best"
- **LF-3-Q2** (5★): "I loved it!"
- **LF-3-Q3** (5★): "I was wonderful session"
- **LF-3-Q4** (5★): "It was really good"
- **LF-3-Q5** (5★): "Excellent"
- **LF-3-Q6** (5★): "Good many tools explored and learned now I'm trying one by one few tried"
- **LF-3-Q7** (5★): "I really enjoyed the session by Larissa Schwartz, where she introduced us to the latest AI image generation tools like Midjourney, Microsoft Designer, DALL · E, Canva, Adobe Firefly, and NightCafe. She…" (response truncated in source)
- **LF-3-Q8** (4★): "Very interesting."
- **LF-3-Q9** (4★): "I think it was a good discussion regarding how to use the different AI image generation tools. A course work based on some of the neural networks behind them

could be a great one."

LF-3.R01 through LF-3.R29 · Individual Luma feedback responses 1 through 29.

LR-3 · Luma roster. 411 registered over the 8-month registration window; 129 attended live; 65% students; 70% expressed interest in an AI Master's program.

E.5 Iteration 4 · GenAI Works Cohort · Sept 2025

The fourth iteration was the second compressed cohort, delivered September 8 through September 12, 2025, through partnership with GenAI Works and broadcast on the GenAI Works YouTube channel. This is the second phase of distillation in my pioneer practice and the iteration with the most extensive recorded teaching delivery in the corpus. The Iteration 4 corpus is the strongest single source for the hallucination-as-pedagogy finding (§D.2) and for the multi-channel teaching practice finding (§D.4) at scale.

E.5.1 Curricular artifacts

The curricular corpus for Iteration 4 is the same workshop deck template as Iteration 3, with date headers updated to reflect the September schedule. The template stability between cohorts is itself evidence for the compression-as-curriculum-maturation finding.

DK-4 · Workshop deck, 22 slides; same template as DK-3 with date headers updated to September 8 through 12, 2025.

E.5.2 Audience and feedback

The Iteration 4 audience data is substantially larger than Iteration 3's, reflecting the GenAI Works partnership's broader reach. The Iteration 4 Luma event page (https://luma.com/m31ygqao) accumulated registrations across the cohort's promotion period, and the YouTube live broadcasts captured real-time participation numbers per session.

LR-4 · Iteration 4 Luma roster. Approximately 4,731 registered guests on the Luma event page across the cohort's registration window. The Iteration 4 partnership with GenAI Works (a

community of approximately six million LinkedIn followers per Day 1 transcript reference) drove substantially higher registration than the CEAS-sponsored Cohort 1.

LR-4.D1 · Iteration 4 Day 1 live participants. 2,654 online participants joined the Day 1 YouTube broadcast on September 8, 2025.

LF-4 · Iteration 4 Luma feedback aggregate. 256 ratings; average 4.2 of 5.

E.5.3 Delivery artifacts

The delivery corpus for Iteration 4 is the largest single text-volume source in the dissertation. The five day-by-day YouTube transcripts together total approximately 55,000 words of contemporaneous teaching delivery, with Day 1 carrying the hallucination teaching passage that supports the most recent anchor for the hallucination-as-pedagogy finding.

TR-4.D1 · Day 1 transcript: Image generation. 10,962 words. Tom Yeh guest from CU Boulder; real-time global audience visible (Nigeria, UK, Denver, Costa Rica named in the first 30 lines).

- **TR-4.D1-Q1**: Opening session, host welcoming attendees and naming the geographies in the chat ("Amazing. Welcome everybody… We got um people from all over the world. I see Nigeria, Denver, UK, Costa Rica.")
- **TR-4.D1-Q2**: Host introducing the workshop's co-hosts ("So, uh we have Tom Yei. welcome from uh you know University of Colorado uh Boulder. We have Lissa Schwarz as well is going to be our main host today.")
- **TR-4.D1-Q3**: Hallucination teaching passage during Microsoft Designer walkthrough ("And you can also see these different hallucinations that are going on. So there aren't even bodies in these shoes… You can keep inputting and inputting and inputting the prompts…")

TR-4.D2 · Day 2 transcript: Video generation. 11,395 words.

TR-4.D3 · Day 3 transcript: Audio generation. 11,599 words.

TR-4.D4 · Day 4 transcript: Research tools. 9,921 words.

TR-4.D5 · Day 5 transcript: Human-Centered AI and Vibe Coding. 11,276 words.

E.6 Cross-iteration evidence · 2023 to 2026

The cross-iteration corpus contains the artifacts that span more than one iteration or sit outside the iteration framework entirely. They include the institutional and program context that situates the entire research arc, my public-facing reflective writing across multiple channels, K-12 outreach work, external media coverage, and federal-research engagement. The cross-iteration evidence is what makes the multi-channel teaching practice finding (§D.4) documentable: it shows the four-theme architecture operating across eight delivery channels and three distinct audience age ranges over more than two and a half years.

E.6.1 Program and grant context

The program and grant artifacts situate the work within its institutional setting. The AI-IRT Seed Grant proposal documents the seeded research arc that connects to both committee members (Tom Yeh and Diane Sieber as co-PIs), and the ENED Preliminary Exam Part 2 records my commitment to offering a continuous generative-AI course at CU Boulder.

AI-PROPOSAL · AI-IRT Seed Grant proposal. Principal investigators Tom Yeh (Computer Science) and Diane Sieber (Herbst Program), both on my committee. Explicit K-12, undergraduate, and graduate scope.

PR-PART2 · ENED Preliminary Exam Part 2, Innovation Proposal. Documents my intent to offer a continuous Generative AI class for CU Boulder because there had been only one such class offered previously (the one I had taught).

E.6.2 Reflective writing

The cross-iteration reflective writing carries my first-person voice across the 2024-2026 period and is the principal supplementary source for the analytic-reflexivity criterion in Iterations 2 through 4 (where the structured weekly journaling of Iteration 1 has no direct counterpart). The Research Impact Essay anchors my practitioner-pioneer biography; the Keep Up Newsletter publishes monthly tool-focused reflections; the Keep Up Podcast extends the same content into audio

form; together they provide the public-facing reflective channel that complements the iteration-internal data.

RE · Research Impact Essay, c. Spring 2024. Three-paragraph self-narrative written in the first person, likely for a fellowship application.

- **RE-Q1**: "Over a year ago I started researching Generative AI and how teachers and students can use different art applications, such as DALL E, Midjourney and NightCafe in order to create artwork from a prompt."
- **RE-Q2**: "I shared my ideas with a school district in Colorado, where I led a professional development to teach the staff how to incorporate and collaborate with these tools in the classroom."
- **RE-Q3**: "From this art contest, we had over 60 students compete and it was such a success that it motivated me to pilot the first ever Generative AI class at the University of Colorado Boulder through the Atlas Institute."
- **RE-Q4**: "During this spring semester, I've been teaching Generative AI to undergraduate students about how these different applications can be used in Education, Industry and Accessibility, and the ethical concerns that arise while using these programs."

KN-EP1 · Keep Up Newsletter Episode 1, Image Generation (April 24, 2025). Public LinkedIn Pulse newsletter; first episode using the running-and-training metaphor.

- **KN-EP1-Q1**: "Expect variable results, occasional hallucinations; persistence improves prompting skills; join community groups" (Lessons-from-Training section).

KN-EP2 · Keep Up Newsletter Episode 2, Research Tools (May 7, 2025).

KN-EP3 · Keep Up Newsletter Episode 3, Sound Tools (May 23, 2025). Publicly credits an Iteration 1 student for introducing Soundful into the curriculum.

- **KN-EP3-Q1**: "I learned about Soundful from one of my students during class." (the Ethan Cuenca → Soundful student-to-instructor tool flow)

KP-EP2 · Keep Up Podcast Episode 2 transcript, May 2025. ~30 minutes; AI by Hand YouTube channel.

KP-EP3 · Keep Up Podcast Episode 3 transcript, May 2025.

E.6.3 HCI summer 2024 guest series

The HCI guest series consists of ten lectures I delivered for a Human-Computer Interaction course at CU Boulder in summer 2024. The series is the most direct transport of the four-theme curriculum architecture into a different course taught by a different instructor; it is the primary evidence for the curriculum-portability claim (§A.5) and contributes to the multi-channel teaching practice finding.

HC-INTRO · HCI Generative AI Intro and Prompt Engineering deck.

HC-ACCESS · HCI Generative AI within Accessibility.

HC-AUDIO · HCI Generative AI within Audio (Ethics).

HC-EDU · HCI Generative AI within Education.

HC-INDUSTRY · HCI Generative AI within Industry.

HC-MUSIC · HCI Generative AI within Music (Ethics).

HC-VIDEO-ETH · HCI Generative AI within Video (Ethics).

HC-VIDEO · HCI Generative AI within Video.

HC-FUTURE-WHEEL · HCI Prompt Engineering using Future Wheel.

HC-SORA · HCI Key Findings in Relation to Sora.

E.6.4 K-12 outreach and research

The K-12 outreach corpus documents my work with children and teens at the CU Boulder STEAM Festival and through the University of Washington KidsTeam and Youth Advisory Board (YAB) co-design sessions in summer 2024. The STEAM Festival mural was a live workshop activity I led; the UW KidsTeam collaboration was a co-design research project I joined as a participant. Together they provide the K-12 end of the audience age range and contribute the near-independent corroboration for the hallucination-as-pedagogy finding (KT-THEMES-C5).

ST-MURAL · AI Art Mural with 50+ Colorado children at the CU STEAM Festival, 2024. IDC conference submission document.

ST-PHOTO · STEAM Festival workshop photograph showing the mural being built live.

KT-DECK · UW KidsTeam Spencer Project slide deck.

KT-IDEAS · UW KidsTeam 2024 three-day AI Ideas session plan.

KT-COMIC · KidsTeam Comic Boarding worksheets (art and math AI-in-schools scenarios).

KT-NOTES · KidsTeam observational and analytic notes on session videos.

KT-YAB · UW Youth Advisory Board (YAB) Multimodal-AI session deck.

KT-THEMES · KidsTeam Research Questions and surfaced Themes, July 2024. Children's and teens' independent observations on GenAI in schools, surfaced through the UW KidsTeam co-design methodology (which the UW team led; I participated).

- **KT-THEMES-C1** (challenges): GenAI being banned from schools
- **KT-THEMES-C2** (challenges): Cheating
- **KT-THEMES-C3** (challenges): AI as "data fed into a computer" that "doesn't yet know everything"
- **KT-THEMES-C4** (challenges): Lack of emotion and human connection
- **KT-THEMES-C5** (challenges): "Hallucinations (such as images produced with a third arm)"
- **KT-THEMES-O1** (opportunities): Teachers allowing GenAI for paragraphs with mistakes to correct
- **KT-THEMES-O2** (opportunities): Checking math homework
- **KT-THEMES-O3** (opportunities): Producing faster work

E.6.5 External media coverage

External media coverage documents how journalists framed the work for non-academic audiences. The single article in the corpus, from Aspen Public Radio, appeared one week after Iteration

1 ended and provides external corroboration that the iteration landed pedagogically.

AP-2024-05-16 · Aspen Public Radio article, May 16, 2024. Title: "Could AI be the next college teaching assistant? Some Colorado professors believe so." Larissa named and quoted; Iteration 1 student Ashley Stafford also quoted. Cited at source-level.

E.6.6 Federal research engagement

The federal-research engagement is the most recent public-facing synthesis of the work, delivered approximately ten weeks before defense. The CU RMACC webinar, uploaded by the federal NAIRR Pilot program, makes the four-theme architecture and the pioneer practice visible to a federal-research audience.

WB-2026-03-03 · CU RMACC Webinar via the federal NAIRR Pilot platform, March 3, 2026. "Unleashing Creativity with Generative AI"; 56 minutes; ~9,600 words of transcribed delivery. Cited at source-level throughout this dissertation.

E.6.7 Other public-facing artifacts

A small number of additional public-facing artifacts complete the cross-iteration corpus. They document teaching activities that did not fit into the main iteration framework but contribute to the multi-channel teaching practice claim.

DK-DLS · DLS Prompt Engineering presentation. Undated; 17 slides; four named participant outputs. Delivered at the CU Discovery Learning Center.

STC · Storytelling Cartoonimator worksheet. K-12-oriented worksheet pairing storytelling with the Cartoonimator tool.

E.7 Conventions

- **Verbatim quotes** preserve original punctuation. Minor clarifications use `[sic]` only when needed for reader comprehension.
- **Q-IDs are assigned on demand** as chapters cite. Adding a quote that does not yet have an ID is a normal step during drafting.

- **New artifacts** added after 2026-05-12 receive IDs at ingestion. The category prefix table above is extensible.
- **Multiple citations in one parenthesis**: separate with semicolons, e.g., `(WU-1.W01-Q1; KT-THEMES-C5; KN-EP1-Q1)`.

Appendix F

AI Use Disclosure

F.1 AI Use Disclosure

This appendix discloses the use of generative-AI assistants during the preparation of this dissertation, in accordance with the transparency expectations that the dissertation itself argues for in its treatment of GenAI pedagogy.

F.1.1 Statement of Authorial Responsibility

I am the sole author of this dissertation. Generative-AI assistants helped with specific tasks of drafting, revising, and structuring described below. All substantive analytic claims, methodological decisions, theoretical findings, and interpretive moves are my own. The conclusions, contributions, and limitations are mine to defend.

F.1.2 AI Tools Used in Dissertation Preparation

Two distinct categories of AI tools relate to this dissertation. Each is named here for clarity.

Tools used as writing assistants during dissertation preparation:

- **OpenAI ChatGPT** used periodically during the spring 2026 preparation period for exploring alternative phrasings and for sanity-checking specific paragraphs.

Tools that are the subject matter of the dissertation:

The generative-AI tools that appear throughout the four iterations as instructional content — ChatGPT, Microsoft Copilot, DALL-E, Midjourney, Stable Diffusion, Runway, Pika, Sora, Soundful, Suno, Udio, ElevenLabs, HeyGen, NotebookLM, Microsoft Designer, and others — were taught and demonstrated in class as part of the curriculum. My own use of these tools to support classroom delivery and curriculum design is documented within the iteration narratives (Appendix C). I do not separately enumerate that classroom usage here.

F.1.3 How AI Assisted with Dissertation Preparation

Concrete enumeration of how AI writing assistants were used during the spring 2026 preparation period:

- **Drafting and revising prose** for both the main dissertation chapters and this appendix, working from outlines, source notes, and substantive content that I provided. Each AI-assisted draft was reviewed by me, often substantially edited, and approved before being incorporated.
- **Restructuring chapter outlines** and section orderings, including the renumbering of the appendix to the letter-prefixed scheme (A through G).
- **Exploring methodological framings**, including the curriculum-as-research framing in §2.6, the autoethnographic-supplement framing of Appendix B, and the reframe of student-facing inputs as course evaluation rather than research data. Each framing was reviewed and approved by my chair and committee before being committed to the document.
- **Bibliography construction**, including the assembly of refs.bib entries with explicit verification flags on uncertain fields (issue numbers, page ranges, DOIs) that I confirmed against the actual publications.
- **Internal-consistency review and claim calibration**, including the audit of student-outcome language that motivated the curriculum-as-research reframe and the softening of contribution claims in §1.5 and §5.1.

F.1.4 What Remained Entirely the Author's

The following were not AI-generated and predate or stand apart from the AI-assisted writing:

- **Empirical artifacts and contemporaneous data** from the four iterations: the Weekly Updates Prelim Document (WU-1.W01–W15) recorded in real time during Iteration 1; the Research Impact Essay (RE); the Keep Up Newsletter (KN-EP series, 2025); the Keep Up Podcast (KP-EP series, May 2025); the CU RMACC webinar transcript (WB-2026-03-03); the five Iteration 4 day-by-day YouTube transcripts (TR-4.D1–D5); the UW KidsTeam co-design corpus (KT-THEMES, KT-DECK, KT-IDEAS, KT-COMIC, KT-NOTES, KT-YAB); the Aspen Public Radio coverage (AP-2024-05-16); the Luma feedback corpora (LF-3, LF-4); and the named student outputs in the DK-1.W01 opening deck. Where I quote any of these, the quotes are verbatim from the source materials, which I authored or collected during teaching long before any dissertation writing began.
- **The four-theme curriculum architecture** (Education, Industry, Ethics, Accessibility) was built into my teaching from my K-12 outreach work onward, before the introduction of GenAI tools. The architecture predates and is independent of any AI-assisted dissertation writing.
- **The three theoretical findings** developed in Appendix D (modularity, learner choice, continuous feedback as the principles, with compression-as-curriculum-maturation, multi-channel teaching practice, and hallucination-as-pedagogy as the autoethnographic sub-claims) are analytic claims I made about my own practice from within that practice. AI assistants helped articulate and refine the prose that explains these claims; the claims themselves are mine.
- **Final decisions** on what to claim and what to hedge, what to retain and what to remove, what to acknowledge as a limitation, and how to characterize the contributions of this work.

F.1.5 Disclosure as Consistent with the Dissertation's Posture Toward AI

The dissertation studies pioneering generative-AI pedagogy. My use of AI assistants in writing about that pedagogy is internally consistent with the curriculum's posture rather than in tension with it. The pedagogical stance documented across the four iterations — name the tool, show what it does and does not do, model responsible use, iterate when the output is inadequate — is the same stance taken in this disclosure: I name the tools used, describe what they did and did not do in this dissertation, and place the human author's judgment as the final arbiter of substantive claims. Anderson's analytic-autoethnography framework licenses reflexive disclosure about the conditions of analytic work, and this disclosure is part of the performed reflexivity that Appendix B §B.7 develops.

F.1.6 Compliance with University Policy

This disclosure is provided in accordance with the University of Colorado Boulder Graduate School's expectations on the transparent use of generative-AI tools in graduate scholarly work, and with the College of Engineering and Applied Science's guidance on responsible AI use in academic writing. The committee was kept apprised of the AI-assisted preparation process throughout. Any university or program policy revisions that arrive between the defense and final submission will be reflected in the final version of this appendix.

Appendix G

References

G.1 References

Adler, P. A., & Adler, P. (1987). **Membership roles in field research**. Qualitative Research Methods Series. Sage Publications.

Anderson, L. (2006). Analytic autoethnography. **Journal of Contemporary Ethnography**, 35(4).

Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. **International Journal of Artificial Intelligence in Education**, 32(4), 1052–1092.

Basgen, B. (2023). **A generative AI primer**. EDUCAUSE Review (Online).

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In **Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)**. ACM.

Bommasani, R., Hudson, D. A., Adjeroh, D., Altman, R., Arora, S., von Arx, S., … Zhang, Y. (2021). **On the opportunities and risks of foundation models**. arXiv preprint arXiv:2108.07258.

Boyer, E. L. (1990). **Scholarship reconsidered: Priorities of the professoriate**. Carnegie Foundation for the Advancement of Teaching.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., … Amodei, D. (2020). Language models are few-shot learners. **Advances in Neural Information Processing Systems**, 33, 1877–1901.

Burrows, A. C., Lockwood, M., Borowczak, M., Janak, E., & Barber, B. (2018). An iter-

ative participatory approach to developing an early childhood problem-based STEM curriculum. **Journal of Educational Research**, 111(3), 256–266.

Chang, H. (2008). **Autoethnography as method**. Left Coast Press.

Chang, H., Ngunjiri, F. W., & Hernandez, K.-A. C. (2013). **Collaborative autoethnography**. Left Coast Press.

Christensen, C. M., Grossman, J. H., & Hwang, J. (2009). **The innovator's prescription: A disruptive solution for health care**. McGraw-Hill.

Cochran-Smith, M., & Lytle, S. L. (2009). **Inquiry as stance: Practitioner research for the next generation**. Teachers College Press.

Code of Federal Regulations. (2018). **Federal Policy for the Protection of Human Subjects (Common Rule), 45 CFR 46.102(l)**. U.S. Department of Health and Human Services.

Daniel, B. (2014). Big data and analytics in higher education: Opportunities and challenges. **British Journal of Educational Technology**, 46(5).

Deci, E. L., & Ryan, R. M. (2000). The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. **Psychological Inquiry**, 11(4).

Deci, E. L., & Ryan, R. M. (1985). **Intrinsic motivation and self-determination in human behavior**. Plenum Press.

Dehouche, N., & Dehouche, K. (2023). What's in a text-to-image prompt? The potential of stable diffusion in visual arts education. **Heliyon**, 9(6).

Ellis, C., Adams, T. E., & Bochner, A. P. (2011). Autoethnography: An overview. **Forum: Qualitative Social Research**, 12(1).

U.S. Department of Education. (1974). **Family Educational Rights and Privacy Act (FERPA), 20 U.S.C. § 1232g; 34 CFR Part 99**. United States Code and Code of Federal Regulations.

Finlay, L. (2002). Negotiating the swamp: The opportunity and challenge of reflexivity in research practice. **Qualitative Research**, 2(2), 209–230.

Healey, M., & Jenkins, A. (2009). **Developing undergraduate research and inquiry**. The Higher Education Academy.

Hutchings, P., & Shulman, L. S. (1999). The scholarship of teaching: New elaborations, new developments. **Change: The Magazine of Higher Learning**, 31(5), 10–15.

Kelly, A., & Sullivan, M. (2023). **ChatGPT in higher education: Considerations for academic integrity and student learning**. Journal of Applied Learning and Teaching.

Kemmis, S., McTaggart, R., & Nixon, R. (2014). **The action research planner: Doing critical participatory action research**. Springer.

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In **Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems**. ACM.

Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). **Intelligence unleashed: An argument for AI in education**. UCL Knowledge Lab.

Lungu, M. (2022). The coding manual for qualitative researchers. **American Journal of Qualitative Research**, 6(1), 232–237.

Marcus, G., & Davis, E. (2020). **Rebooting AI: Building artificial intelligence we can trust**. Pantheon Books.

McKenney, S., & Reeves, T. C. (2018). **Conducting educational design research** (2nd ed.). Routledge.

Mishra, P., & Koehler, M. J. (2006). Technological Pedagogical Content Knowledge: A framework for teacher knowledge. **Teachers College Record**, 108(6), 1017–1054.

Mollick, E. R., & Mollick, L. (2023). **Assigning AI: Seven approaches for students, with prompts**. Working paper.

Office for Human Research Protections. (2014). **Guidance on distinguishing quality improvement and research**. U.S. Department of Health and Human Services.

OpenAI. (2023). **GPT-4 technical report**. Retrieved from https://openai.com/research/gpt-4

Prensky, M. (2001). Digital natives, digital immigrants. **On the Horizon**, 9(5), 1–6.

Prince, M. J., & Felder, R. M. (2006). Inductive teaching and learning methods: Definitions, comparisons, and research bases. **Journal of Engineering Education**, 95(2), 123–138.

Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., & Sutskever, I. (2021). Zero-shot text-to-image generation. **Proceedings of the 38th International Conference on Machine Learning**, 8821–8831.

Redecker, C., & Punie, Y. (2017). **European Framework for the Digital Competence of Educators: DigCompEdu**. Publications Office of the European Union.

Reed-Danahay, D. (Ed.). (1997). **Auto/Ethnography: Rewriting the self and the social**. Berg.

Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. **American Psychologist**, 55(1), 68–78.

Ryan, R. M., & Deci, E. L. (2017). **Self-determination theory: Basic psychological needs in motivation, development, and wellness**. Guilford Press.

Salmon, G. (2002). **E-tivities: The key to active online learning**. RoutledgeFalmer.

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

Schunk, D. H. (2020). **Learning Theories: Educational Perspectives** (8th ed.). Pearson.

Selwyn, N. (2021). **Education and technology: Key issues and debates**. Bloomsbury Publishing.

Taleb, N. N. (2007). **The black swan: The impact of the highly improbable**. Random House.

Tedlock, B. (1991). From participant observation to the observation of participation: The emergence of narrative ethnography. **Journal of Anthropological Research**, 47(1).

Tomlinson, C. A. (2014). **The differentiated classroom: Responding to the needs of all learners**. ASCD.

Touretzky, D., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019). Envisioning AI for

K-12: What should every child know about AI? In **Proceedings of the AAAI Conference on Artificial Intelligence**, 33.

UNESCO. (2021). **Artificial intelligence in education: Challenges and opportunities for sustainable development**. United Nations Educational, Scientific and Cultural Organization.

UNESCO. (2021). **Equity, inclusion and learner diversity**. UNESCO Publishing.

Wendell, K. B., Koehler, C. M., Fanucchi, A. K., & Baral, K. (2021). Designing biomimetic robots: Iterative development of an integrated technology design curriculum. **Educational Technology Research and Development**, 69(6), 3141–3164.

Wobbrock, J. O., & Kientz, J. A. (2016). Research contributions in human-computer interaction. **Interactions**, 23(3).

Yin, R. (2014). **Case study research: Design and methods** (5th ed.). Sage.