Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A longitudinal constructionist case study of five student positions on AI as a creative collaborator in open source AI art pedagogy

Atticus Sims

atticussims@um.edu.mo

University of Macau

Research Article

Keywords: generative AI in higher education, constructionist learning, arts pedagogy, longitudinal case study, reflexive thematic analysis, asynchronous email interviewing

Posted Date: July 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9713097/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A longitudinal constructionist case study of five student positions on AI as a creative collaborator in open source AI art pedagogy

Atticus Sims
Department of Arts and Design, Faculty of Arts and Humanities
University of Macau, Taipa, Macau SAR, China
atticussims@um.edu.mo ORCID: 0009-0009-2933-2344

14 May 2026

Abstract

The constructionist tradition predicts that "low floor, wide walls" pedagogical environments produce individually-different student outcomes, with the differentiation itself the pedagogical achievement. Few studies characterise this at the granularity of a structured spectrum on a focal technology, and fewer still with post-graduation longitudinal data. This paper presents a five-subject longitudinal multiple case study of a single cohort across a two-year zemi (the Japanese seminar structure that sustains the same cohort and professor across multiple academic years) at Kyoto Seika University. The pedagogy was constructionist by design; the tool stack evolved from ComfyUI—an open-source node-based interface for generative AI models—in Semester 1 into a mixed ecosystem of AI-assisted code editing, proprietary video generation, large language models, and game engines per project. The cohort was tracked through their March 2026 graduation across a ∼25-month observation window, with a 2024 end-of-instruction qualitative survey and April–May 2026 email interviews forming the two-timepoint analytical pair, triangulated against instructor–student Discord exchanges, Production Notes, and practitioner reflective documentation. Findings document a five-position spectrum of AI-as-creative-collaborator stances—expansive system-designer, AI-as-operations-manual, conceptual-amplifier-only, bounded-ethical-and-pragmatic, and bounded-philosophical. An independently-constructed mental-model evolution typology over the 2024–2026 longitudinal pair converges one-to-one onto the same five-position structure, providing within-cohort triangulation across two analytic paths constructed by the same researcher. An asymmetric-amplification analysis identifies four candidate mechanisms by which AI scaffolding can fall short at the realisation layer, each anchored at one student case and corroborated by that subject's 2024-baseline help-seeking disposition.

Keywords: generative AI in higher education; constructionist learning; arts pedagogy; longitudinal case study; reflexive thematic analysis; asynchronous email interviewing.

1 Introduction

Constructionist learning theory predicts that pedagogical environments designed around what Resnick [1] calls "low floor, wide walls, high ceiling"—building on Papert's earlier articulation of "low floor, high ceiling" in *Mindstorms* [2]—produce individually-different student trajectories, and that this differentiation is itself the pedagogical achievement [1–4]. The prediction is well-established at the level of the *artefacts* students produce—diverse projects, varied creative directions, distinct technical strategies—and extensively supported in the constructionist literature on programming environments, robotics, digital media, and creative coding. It is far less tested at the granularity of a *structured spectrum of intellectual stances* on the focal technology,

and almost untested with post-graduation longitudinal data. This paper asks whether, in an arts-pedagogy environment built around generative AI tools, one constructionist seminar can produce a small number of recognisable, defensibly-different student positions on what AI is for and how it should relate to creative authority—and whether those positions persist in students' post-graduation practice.

Generative AI in higher education has become a rapidly stratified literature in the period between 2023 and 2026 [5–11]. Within that literature, work on creative-arts curricula has converged on a small set of recurrent concerns: how to scaffold student–AI interaction so that AI does not displace student creative authority; how to integrate generative tools into existing studio and critique cultures; and how to assess outcomes when the artefacts students produce are themselves the joint product of student decisions and machine generation. The present paper engages these concerns from one such empirical site, positioning itself as an *AI in higher education* paper using arts pedagogy as the site rather than an *AI art pedagogy* paper for an arts-only readership.

The study is a qualitative multiple case study of a two-year zemi in the Image Expression Major at Kyoto Seika University's Faculty of Media Creation. The zemi is a Japanese higher-education seminar structure with no exact Western equivalent: a single cohort of undergraduates studies with the same supervising professor across two or more academic years, integrating project-based learning with sustained mentorship and a community of practice that runs from third year through graduation [12]; full comparative framing in Section 4.1. The cohort tracked here entered the zemi as third-year undergraduates in April 2024 (nine enrolled, weekly in-class meetings); the same cohort continued into the 2025–2026 fourth-year graduation-supervision phase (advisory format, on-demand supervision). ComfyUI—an open-source node-based interface for generative AI models, used here primarily with Stable Diffusion image generation but extensible to FLUX, video diffusion variants, LoRA fine-tuning, ControlNet conditioning, and community-contributed workflows—served as the primary tool in Semester 1; the curriculum then expanded to AI-assisted code editing (Cursor), proprietary video generation (RunwayML), large language models (Grok, Gemini), and game-engine and node-graph environments (Blender, Unity, Unreal Engine, TouchDesigner) as individual projects required. The pedagogy was articulated as constructionist by the practitioner before the zemi began—in a December 2023 lecture-to-students that named Papert directly and in a December 2024 preprint [13] that articulated the framing in writing; the framework-genealogy claim is developed in Section 2.4. Five students from the cohort are tracked across the two-year arc through their March 2026 graduation exhibition and into early post-graduation practice via asynchronous email interviews conducted in April–May 2026, two months after graduation; the longitudinal observation window spans the full zemi arc ($\sim$25 months from the April 2024 start to the Round 1 interviews).

The paper makes five substantive contributions. First, it documents a *five-position spectrum* of AI-as-creative-collaborator stances populated by five students from a single constructionist zemi—expansive system-designer; AI-as-operations-manual; conceptual-amplifier-only; bounded-ethical-and-pragmatic; bounded-philosophical—at a granularity not previously documented in the AI-art-pedagogy literature. Second, the spectrum is *triangulated within-cohort* by an independently-constructed mental-model evolution typology over the 2024–2026 longitudinal pair: two analytic paths using different data sources, temporal anchors, and coding procedures converge one-to-one onto the same five-position structure. Third, an *asymmetric-amplification* finding with four mechanisms and three workaround paths: AI scaffolds non-technical-student work at the conceptual layer but fails at the realisation layer in mechanism-specific ways, with each mechanism corroborated by its subject's 2024-baseline help-seeking disposition. Fourth, a *multi-modal register-mix mentorship* typology in which uptake operates across multiple modes simultaneously within a single student's case, distributed across mentorship sources, content layers, and supervision-frequency contexts. Fifth, a methodological argument: pairing an end-

of-instruction qualitative-survey baseline with a longitudinal follow-up interview round, triangulated against instructor direct-message records, student-authored Production Notes, and practitioner first-person interpretive content, produces a cross-validation dividend that interview-only or instructor-records-only analysis cannot deliver—and the architectural feature that makes the dividend operate (practitioner memory's structural asymmetry under longitudinal-retrospective design) is itself a transferable resource for practitioner-inquiry case-study research in AI in higher education.

2 Background and theoretical frame

2.1 Constructionism and the prediction of individuated outcomes

Constructionism [2, 3, 14] frames the study. Where Piaget's constructivism describes knowledge as built rather than received, Papert added the further claim that knowledge construction is most powerful when learners are *building shareable, public artefacts*—programs, sculptures, simulations, exhibitions—that externalise their thinking and make it available for revision. Resnick's [1] reformulation of the constructionist environment as a *low floor, wide walls, high ceiling* design space gives the theory a pedagogical operationalisation: low floor for accessible entry regardless of technical background, high ceiling for sophistication that rewards sustained engagement, and wide walls for the diversity of legitimate trajectories the environment must support. Kafai's [4] instructionist-versus-constructionist distinction in game-based learning sits within a broader Kafai-led strand that has extended constructionism into digital media and computational participation [15] and into AI-tool-mediated making more recently [16], is the line of theoretical inheritance the present study draws on most directly. Recent work has continued to extend constructionism into the generative-AI era specifically: Levin, Semenov and Gorsky [17] trace the trajectory of constructionist principles across three digital epochs culminating in generative AI, and Lin, Dai and Ng's [18] systematic review of K-12 AI literacy education documents the field's pedagogical-design patterns at school level. The present study extends this trajectory into the undergraduate arts-pedagogy site.

Three features of the constructionist tradition matter for what the paper claims. First, the tradition makes a substantive prediction about *outcome variance*—that the wide-walls environment will produce individually-different student trajectories, and that this differentiation is itself the pedagogical achievement, not noise around a uniform target. Second, the tradition treats the *constructed artefact* as the unit of evidence: students' projects, code, exhibitions, and self-narrations of process are the materials through which learning is read. Third, the tradition is methodologically hospitable to interpretive case study research; Papert's own writing combines theoretical argument with extended, narrative readings of individual learners, and this analytical move—naming a stance, anchoring it in a single learner's record, and asking what it implies for design—recurs in the present paper's case sketches.

The prediction of individuated outcomes is empirically under-tested at the granularity of a *structured spectrum of stances* on a focal technology, and even less tested with longitudinal data tracking those stances into post-graduation creative practice. Studies of constructionist environments have repeatedly documented that students produce different projects, develop different mental models, and adopt different relationships to the tools used [1, 15]; few studies have analysed those differences as a *spectrum of positions on the focal technology itself*, and fewer still have done so where the technology in question is generative AI. The empirical contribution of the present study lies in this gap.

A subsidiary point: the constructionist framing of "wide walls" is sometimes read as predicting only a diversity of *projects*, not a structured spectrum of intellectual *stances* on the medium. The reading proposed here—that constructionist outcomes include differentiated stances on

what the medium is for and how it should relate to creative authority—is a stronger claim than the standard wide-walls reading and is the analytical extension the paper's findings will support. Resnick's [1] own treatment of the *Four Ps* (projects, passion, peers, play) gestures in this direction without operationalising it; the present paper's contribution is to show that, given the right empirical access, the differentiation can be characterised at the level of stance, not just at the level of artefact.

2.2 Reflective practice and practitioner inquiry

The methodological warrant for the practitioner's reflective documentation as primary data is supplied by Schön's [19, 20] account of *reflective practice* and Cochran-Smith and Lytle's [21, 22] account of *practitioner inquiry*. Schön's distinction between *knowing-in-action* (the tacit competence visible in expert performance) and *reflection-in-action* and *reflection-on-action* (real-time and retrospective examination of practice, respectively) describes the form of professional knowledge the present study treats as evidence. The practitioner's pedagogical decisions—when to introduce a tool, when to delegate technical questions to AI rather than answering them, when to deploy a long-form direction letter, when to redirect a student's medium choice—are not derivable from educational theory alone; they are instances of what Schön [20] called *professional artistry*, expert judgment under conditions of uncertainty, uniqueness, and value conflict, and Polanyi's [23, 24] earlier argument that knowing always exceeds saying provides the broader epistemological underpinning.

Cochran-Smith and Lytle's [21] tripartite distinction between *knowledge-for-practice* (formal theory applied), *knowledge-in-practice* (craft knowledge embedded in expert performance), and *knowledge-of-practice* (knowledge generated when practitioners treat their own work as a site of inquiry) places the present study in the third category. Their [22] framing of practitioner inquiry as an *inquiry-as-stance*—an ongoing professional disposition rather than a method deployed temporarily—is the orientation the paper assumes. Practice-led research traditions in art and design [25–27] provide an adjacent warrant for treating creative-practitioner expertise as research-generative, and are cited here in a light register; the present paper is not principally a practice-led-research paper but an arts-pedagogy paper that draws on the practitioner-inquiry tradition for its methodological footing.

2.3 Generative AI in higher education

The 2023–2026 literature on generative AI in higher education has expanded rapidly and stratified across disciplines. Within the present paper's empirical neighbourhood, three lines of work are most relevant. First, studies of generative-AI integration into creative-arts curricula [5, 7–9, 28, 29] document the early period of curricular experimentation in which the present zemi sits; Hwang and Wu's [9] TPACK-framed account of text-to-image studio integration in undergraduate graphic design is the closest art-education-journal analogue to the present empirical site. Second, work that extends generative-AI inquiry beyond higher-education-arts into adjacent levels—K-9 socio-material studies of student–AI co-creation [30, 31] and HE-arts constructivist-learning studies [6]—frames the cross-level pedagogical questions the present study engages. Third, two recent systematic reviews of generative AI in art education—Jiang, Fan and Liu [10] and Fang [11]—map the current field and triangulate each other across overlapping empirical corpora; they are used here as reference points for situating the paper's empirical contribution.

A through-line worth naming is the open-source-versus-proprietary distinction, which has consequences both for what students can inspect and modify and for the economic accessibility of the tools to which student creative practice becomes coupled. The seminar studied here used ComfyUI—an open-source node-based interface for a wide range of generative AI models

(Stable Diffusion variants, FLUX, video diffusion models, LoRA fine-tuning, ControlNet conditioning, audio-generation pipelines, and community-contributed workflows) in which every step of the generative pipeline is visible, parameterisable, and modifiable—as the *constructionist transparent-pipeline anchor* in Semester 1 of Year 1, with the tool stack expanding subsequently to include AI-assisted code editing, proprietary video generation, large language models, and game-engine environments per project (described in Section 4.2). This breadth is itself part of the constructionist-transparency argument: a node-graph interface that generalises across model architectures supports deeper mental-model formation than a model-specific UI, because the student learns the operational paradigm rather than a tool-specific surface. The pedagogical case for open-source tools at the constructionist-anchor point rests on a *transparency* argument: tools whose internal processes are inspectable function as Papert's [2] *objects-to-think-with*, supporting deep mental-model formation in ways that opaque, prompt-only commercial services do not. The present paper takes this open-source framing seriously at the anchor point without making the stronger claim that constructionist outcomes are achievable *only* with open-source tools; the seminar's own subsequent integration of proprietary services alongside the open-source anchor—in the broader context of the rapid 2023–2025 emergence of commercially-available text-to-image generation tools (DALL-E, Midjourney, Stable Diffusion variants) that Tatar et al. [7] characterise as a paradigm shift in artistic practice, and that Zhou, Zhu, Mateas and Wardrip-Fruin [32] examine through their analysis of what text-to-image models offer for visual creativity—suggests that the question is more nuanced than the open-source-versus-proprietary binary alone allows. The paper's findings on mental-model differentiation in Section 6.4 supply some empirical anchoring for that nuance.

Two strands of recent empirical and theoretical work frame the analytical apparatus the present paper deploys in Section 6. The first concerns the documented tension between individual creativity gains and collective output homogenisation under generative-AI use: Doshi and Hauser [33], in a 293-writer experiment using GPT-4, demonstrated that access to generative AI raised individual creativity ratings—concentrated among less-creative writers, producing a levelling effect on individual output quality—while making the resulting outputs measurably more similar to one another, indicating reduced collective novelty. The five-position spectrum documented at Section 6.1 of the present paper supplies one anchored counter-instance to that homogenisation prediction at the level of student stance: under a sustained constructionist zemi design with multi-modal mentorship, the same focal technology produced systematically *differentiated* student trajectories rather than homogenised outputs. The second strand concerns the cognitive-psychological mechanisms by which AI scaffolding succeeds or fails. Risko and Gilbert's [34] framework of *cognitive offloading*—the use of external supports to alter the information-processing requirements of a task—argues that propensity to offload is shaped by metacognitive evaluation of one's own abilities, which can be miscalibrated and produce suboptimal offloading decisions. Crowston and Bolici [35] extend the framework to AI-system contexts specifically, distinguishing the *levelling effect* (AI bringing novices toward expert output without their developing the underlying skill) from the *multiplier effect* (AI amplifying existing expert capability), with prompting, evaluation, and editing as the loci where new competencies are or are not developed. The Section 6.3 asymmetric-amplification mechanisms documented here are read against both frameworks: as miscalibrated-offloading patterns at the realisation layer (Risko and Gilbert), and as a partial mapping of stance-by-mechanism interactions onto the levelling-versus-multiplier distinction (Crowston and Bolici). Recent empirical work on AI integration in undergraduate design and arts coursework documents a mixed picture across outcome constructs and learner profiles (Chen et al. [36]; Sarıtepeci and Yıldız Durak [37]). Wang, Zhang and Kang [38] document *differential* effects on creative self-efficacy, engagement, and design performance moderated by prior performance level and AI tool choice—the closest contemporary anchor for the prior-ability and stance heterogeneity claims developed in Section 6.

The literature on AI in higher education that treats arts pedagogy as one site among many has

converged on a small set of recurrent concerns: how to scaffold student–AI interaction so that the AI does not displace the student's creative authority [7, 30]; how to integrate generative tools into existing studio and critique cultures without flattening their pedagogical specificity [8, 28]; and how to assess outcomes when the artefacts students produce are themselves the joint product of student decisions and machine generation [6]. The present paper engages these concerns from the empirical site of one undergraduate arts seminar; its contribution is not a system-design or assessment-rubric proposal but a longitudinal characterisation of the differentiated stances students take up when a constructionist studio pedagogy is sustained around generative-AI tools and an instructor's expert mentorship operates across multiple registers simultaneously. Jiang, Fan and Liu's [10] systematic review of AI in art education and Fang's [11] near-simultaneous review of the same field together provide the closest field-defining maps to date and are used in Section 7 as reference points for situating the paper's findings within the broader literature.

2.4 Framework-genealogy: the pedagogy was constructionist by design

One framework-genealogy claim underwrites the analysis: the pedagogy was articulated as constructionist by the practitioner *before* the seminar began, not theorised as constructionist afterwards. Two pieces of primary-source evidence support the claim. First, a December 2023 lecture-to-students delivered at Kyoto Seika University and preserved in the practitioner's institutional workspace stated the framing directly: "Our aim is not to focus on teaching technical skills for specific tasks. Instead, we emphasize learning how to acquire new skills in immersive media, aligning with Seymour Papert's Constructionism philosophy, which advocates for 'learning by making'." Second, the practitioner's reflection on the 2023 workshop that preceded the 2024 zemi was published as a non-peer-reviewed preprint in December 2024 [13] and articulated the pedagogy as constructionist in §3.5 of that document—seventeen months before the cohort's graduation exhibition and at a point when the longitudinal data analysed here did not yet exist. The paper's reading of the seminar as a constructionist environment is therefore not a post-hoc theoretical overlay; it is the framing the practitioner deployed in advance, made publicly available, and tested against a class of students who completed the seminar under that framing.

3 Methods

The study reported here is a qualitative multiple case study [39] of a single undergraduate seminar on AI art at a Japanese art university, conducted in 2024 and tracked longitudinally into the cohort's post-graduation creative practice through April 2026. The unit of analysis is the seminar; the embedded units of analysis are the five longitudinal student subjects. A sixth subject from a 2023 workshop is treated as a supplementary transferability anchor rather than as a sixth embedded unit. The design integrates *five* streams of qualitative evidence—student interviews, instructor–student direct-message records, student-authored Production Notes, practitioner reflective documentation, and a 2024 end-of-course qualitative survey (1C) that anchors the longitudinal baseline of the two-timepoint pair—and analyses them through reflexive thematic analysis [40–42].

3.1 Research design

The study uses Yin's multiple case study with embedded units of analysis [39], where the case is the 2024–2026 two-year zemi in the Faculty of Media Creation at Kyoto Seika University and the embedded units are five students from that cohort tracked into their March 2026 graduation exhibition and beyond. The case selection is purposive but constructed: the zemi is the primary

site at which the practitioner's constructionist AI-art pedagogy was implemented at scale, and the five subjects span the cohort's full range of AI engagement, from one student who built a tightly-coupled human–AI generative pipeline as her graduation work to one who deliberately constrained AI to technical-research support and made the central creative decisions without it. An additional Round 1 interview was conducted in parallel with a participant from the practitioner's 2023 university-wide ComfyUI workshop (now an active professional generative-AI artist); his returns inform Section 6.5 as a workshop-format comparison reference but are not analytically central to the focal case.

The design is qualitative-led. No statistical generalisation or inferential tests are reported. The analytical strategy is analytic generalisation in Yin's [39] sense: findings are extended to the theoretical propositions of constructionist learning theory and practitioner inquiry, not to a population of students or seminars.

The pedagogical iteration across the practitioner's teaching career is framed as *organic expert iteration*: semester-to-semester refinement grounded in professional artistry [19, 20] and ongoing practitioner inquiry [21, 22].

3.2 Participants and recruitment

The zemi cohort enrolled nine third-year undergraduates in the Image Expression Major at Kyoto Seika University in April 2024. Eight returned full end-of-course survey responses at the close of Semester 1; one was a low-engagement outlier and is not part of the longitudinal corpus. Five engaged students are tracked here as embedded units, pseudonymised as XJ, LX, LY, LJ, and WS. The 2023-workshop comparison participant (pseudonymised as TZ) is not an embedded unit; his returns are used in Section 6.5 only.

Recruitment for the longitudinal interview corpus was conducted in April 2026, approximately one month after the cohort's March 2026 graduation exhibition; longitudinal observation began at the seminar's April 2024 start and the interview-round date is at the end of a ∼25-month observation window across the two-year seminar arc plus two months of post-graduation practice. Recruitment messages were sent individually and explicitly invited each student's participation regardless of how much they had used AI in their subsequent work. The recruitment of LY in particular is documented verbatim—"I know that you didn't rely on AI a lot for your work, but I think that using your work as a case study will be excellent for research"—as evidence that low-AI-use trajectories were sought as part of the design, not retroactively framed as negative cases. Informed consent was obtained in each interview thread; participants were told they could decline any question, withdraw at any time, and that participation had no bearing on past or future academic relationships.

3.3 Data sources and triangulation

Five streams of evidence underpin the analysis. *(1) Round 1 asynchronous email interviews* with the five longitudinal subjects (April–May 2026) used a 12-question semi-structured protocol organised in five parts (creative work, learning ecology, instructor's role, understanding and tools, reflection); a parallel 7-question longitudinal protocol covering graduation-work follow-up was administered in the same exchanges. *(2) Round 2 follow-ups* (three of five returned at draft date) used 5–7 targeted depth-, cross-case-, clarification-, and interpretive-validation prompts derived from cross-case patterns surfacing in Round 1 [43]. *(3) Per-student Production Notes* (Seisaku Nōto, 9–12 pp.) are student-authored submissions written for the Faculty of Media Creation's oral examination at the close of the graduation year, interleaving conceptual framing, prior-research positioning, technical-validation documentation, and significance statements at an academic register more typical of postgraduate than of undergraduate art-school

documentation. As primary-source data they are unusual—independently composed by the student, produced for an institutional examination rather than for the present study—and they function as member-generated material that triangulates against both interview self-report and instructor observation. They are available for four of the five longitudinal subjects (XJ, LX, LY, WS). Production Notes were authored by students for institutional oral examination and may have been refined under supervisor feedback during their composition; their epistemic status is therefore as member-authored, institutionally-evaluated documents rather than research-instrument-generated data—valuable for academic register and graduation-work-formulation temporal anchoring, but not analytically independent of practitioner influence. *(4) Practitioner reflective documentation* covers two distinct but related sub-corpora treated as a single stream throughout the paper: ten teaching documents authored across 2023–2025, and per-student supervision-moment records reconstructed from approximately 1,275 substantive instructor–student exchanges across five Discord channels (October 2023–April 2026), including four long-form direction letters from a May–June 2025 cluster quoted from redacted-copy versions of the originals.

(5) The 2024 Seika 1C end-of-course survey (Instrument 1C: Seika Zemi End-of-Course Survey, ComfyUI-specific) is the fifth primary stream, administered on 28 October 2024 immediately after the Semester 1 ComfyUI-instruction block ($n = 9$ enrolled; 8 complete responses). The 1C instrument captured 17 Likert items on technical experience, learning experience, and working patterns, plus 10 open-ended qualitative items on understanding-change, unexpected discoveries, mental-model metaphors, and help-seeking dispositions. *The 1C → 2026-interview pair constitutes the two-timepoint longitudinal architecture of the case study*: the 1C captured students' articulations at end-of-instruction; the 2026 Round 1 and Round 2 interviews captured the same students' articulations approximately eighteen months later. Quantitative Likert items from the 1C are reported descriptively where they contextualise an interview claim and are not the basis of inferential claims (consistent with the analytic stance described in Section 3.1); the qualitative items are the analytic anchor for the longitudinal-baseline-to-endpoint comparison developed in Section 5 (case sketches) and Section 6.4 (mental-model evolution arcs). Two temporal windows are referred to throughout: the *observation window* (∼25 months, April 2024 zemi start to April–May 2026 Round 1 interviews) and the *longitudinal two-timepoint interval* (∼18 months, October 2024 1C survey to April–May 2026 Round 1).

A parallel cohort of postgraduate students at the University of Macau received a closely-aligned protocol in April–May 2026; their returns will inform a future cross-case extension and are not analysed here.

Triangulation is structural, not a generic appeal to converging evidence. Each substantive finding reported in Section 6 is anchored at the per-student level by at least two of the five streams; cross-case claims are anchored by three or more, with two findings in particular (Section 6.4 mental-model evolution; Section 6.3 asymmetric-amplification mechanisms) anchored by at least three streams, with the 2024 baseline as an additional corroborative path beyond the 2026-endpoint sources. The triangulation strategy is itself a finding, addressed in Section 6.6 and revisited in Section 3.6: practitioner memory under longitudinal-retrospective design is event-anchored at register-escalation moments and post-zemi attenuated-channel conversations, and thinnest at register-baseline supervisory rhythm and student-low-disclosure stretches; the DM transcript, Production Notes, 1C baseline, and interview returns triangulate the stretches that practitioner memory does not retain.

3.4 Asynchronous email interviewing as design choice

Asynchronous email interviewing was chosen for substantive methodological reasons [43–45]: deliberative, pre-transcribed written responses; cross-linguistic flexibility (Japanese / Chinese

passages where participants felt they could express ideas more precisely); and Dahlin's [43] explorative strategy in which first-round patterns inform targeted second-round follow-ups. The format produced *interpretive friction*: across Round 1 and Round 2, most of the five longitudinal subjects issued first-person corrections or refinements of researcher pre-readings; three representative examples follow. WS corrected a pre-interview framing of his advancement-exhibition piece as a Refik-Anadol-inspired particle work pivoted away from; in his account, the particle direction was *abandoned*, both pieces carrying a consistent 3D-window aesthetic. LY refined the pre-interview craft-defensive hypothesis into a philosophically grounded position on embodiment, Round 2 re-anchoring verbatim: "Artificial intelligence is undeniably useful within defined parameters, yet the core of its creativity remains rooted in embodied human experience. I believe this to be a fundamental truth." LX deepened his rationale across rounds (bounded-epistemic → bounded-ethical-and-pragmatic; Section 5.2 / Section 6.2). These corrections are treated as suggestive evidence that the format is methodologically generative. Trade-offs—forfeit of non-verbal cues, spontaneous reactions, synchronous flow—are acknowledged, but for creative-process, technical-understanding, and reflective-evaluation questions the deliberative quality of written responses is an analytical asset.

3.5 Analytic approach: reflexive thematic analysis

Analysis follows Braun and Clarke's six-phase reflexive thematic analysis [40–42]: themes are constructed through interpretive engagement rather than discovered, and researcher subjectivity is an analytical resource, not a contaminant. The six phases—familiarisation, generating initial codes, constructing initial themes, reviewing themes against the dataset, defining and naming themes, producing the analytic narrative—were worked through within-case before cross-case synthesis began. Within-case familiarisation proceeded student-by-student: the analyst read all five streams for one subject in a single block before turning to the next, generated initial codes against that within-case material, and stabilised a per-student interpretive narrative before any cross-case comparison was attempted. This sequencing made it harder for an early cross-case pattern to overwrite a within-case nuance.

Consistent with the reflexive thematic analysis tradition [41], the analysis is single-researcher and treats researcher subjectivity as an analytical resource rather than as measurement noise; trustworthiness is established through the transparent chain of evidence linking each claim to a primary-source anchor and through the calibrated reflexivity disclosed in Section 3.6. LLM use in the research process is disclosed in the Declarations section.

3.6 Reflexivity

The researcher is the practitioner of record for the seminar studied here. Prominent disclosure is consistent with reflexive thematic analysis [41] and practitioner inquiry [22]. Two structural conditions on the data warrant joint disclosure: the August 2025 institutional transition from Kyoto Seika to the University of Macau coincided with a *supervisory regime change* (Phase-1 biweekly critique through July 2025 → Phase-2 on-demand from August 2025) under which the principal development of all five graduation projects occurred, introducing per-student documentary asymmetry; and *practitioner memory under longitudinal-retrospective design is structurally asymmetric*, event-anchored at register-escalation moments and post-zemi attenuated-channel conversations and thinnest at register-baseline supervisory rhythm and student-low-disclosure stretches. Triangulation across the five data streams was specifically configured to address both conditions.

Reflexive practice in this study is calibrated rather than performative: bias-disclaimers and bias-ownerships are stated as the evidence supports (see Section 5.4 LJ supervision-frequency

disclaimer; Section 7.2 LY biographical-affinity ownership) rather than displaced into reflexive throat-clearing. The paper terms this *calibrated reflexivity* and treats it as methodologically more defensible than universal reflexive disclosure in practitioner-inquiry case study research.

3.7 Ethics

The study was approved by the Kyoto Seika University Ethics Review Committee, and the research was conducted in accordance with the Kyoto Seika University Research Ethics Regulations and the Kyoto Seika University Guidelines for Research Involving Human Subjects. The research activities reported here, including the longitudinal follow-up interviews conducted in April–May 2026 and the use of student creative works for research purposes, were supported by the President's Designated Research Grant (*Gakuchō-Shitei-Kenkyū*, 2024). All longitudinal subjects provided informed consent for interview content, Production Note excerpting, and verbatim quotation of supervision records. Pseudonyms are applied throughout; redacted-copy versions of Production Notes and instructor records are used for any verbatim passage; identifying details (real names in any script, student IDs, photographic location-naming filenames) are stripped from all material that may be reproduced or appendiced.

4 Seminar context

4.1 The Japanese zemi as a pedagogical structure

The Japanese *zemi* has no exact Western equivalent and warrants brief comparative framing for readers unfamiliar with its pedagogical scope [12, 46]. Where a Western undergraduate seminar typically runs for a single semester or term with a fresh cohort each iteration and a primarily academic focus on thematic debate and analysis, a zemi runs across two or more academic years with a stable peer cohort and the same supervising professor throughout. It integrates project-based learning with pastoral care and character development; it includes social functions (picnics, field trips, year-end parties) that build a long-term community of practice; and it typically culminates in a graduation thesis (*sotsugyō ronbun*) or large-scale graduation project. For Western readers, the zemi is best understood as a *liminal space* mediating between academic study and the transition to professional creative life, in which students operate as apprentices under a sustained mentor relationship rather than as participants in a discrete academic offering. The two-year zemi arc described in this paper is therefore not "two seminars" in the Western sense but a continuous mentor-student community of practice extended across the full graduation arc, with the August 2025 institutional transition introducing a regime change discussed in Section 3.6 and revisited in Section 4.2 below.

4.2 The 2024–2026 zemi arc

The seminar is a two-year zemi arc in the Image Expression Major at Kyoto Seika University's Faculty of Media Creation, running April 2024 through the March 2026 graduation. The arc comprises two formally distinct phases. *Year 1—the 2024 third-year zemi* (April–November 2024 in its formally scheduled phase, with continued informal engagement through Year-2 transition) enrolled nine third-year undergraduates and operated on a weekly in-class meeting cadence with close instructor oversight. *Year 2—the 2025–2026 fourth-year graduation-supervision zemi* (April 2025–March 2026) operated in an advisory format: biweekly critique meetings through July 2025 (Phase 1), shifting to on-demand-only supervision from August 2025 onwards (Phase 2) following the practitioner's institutional transition from Kyoto Seika to the University of Macau. Of the nine students enrolled in Year 1, eight engaged substantively

across the full two-year arc; one is documented as a low-engagement outlier and is not part of the longitudinal corpus reported here. The principal development of all five longitudinal subjects' graduation projects occurred under Year-2 conditions.

The seminar's tool ecosystem evolved across the two years. *In Semester 1 of Year 1* (∼April–July 2024) ComfyUI was the primary tool, introduced as the constructionist transparent-pipeline anchor around which the first-half curriculum was organised—an open-source node-based interface for a wide range of generative AI models (Stable Diffusion variants, FLUX, video diffusion, LoRA fine-tuning, ControlNet conditioning, community workflows). *After Semester 1 the tool stack broadened*: LLMs (Grok, Gemini, ChatGPT, Claude) for prompt refinement and ideation; AI-assisted code editing (Cursor); proprietary video services (RunwayML and others); and game-engine and node-graph environments (Blender, Unity, Unreal Engine, TouchDesigner) adopted per project. ComfyUI remained available throughout but was the primary tool only for Semester 1; the trajectory is one of expansion from a single open-source constructionist anchor into a mixed ecosystem aligned with each student's developing project requirements. Year 1's delivery used a quarterly cadence (1Q–4Q) with a parallel commitment in the Faculty's Social Practice (*shakai jissen*) practicum on Kyoto cultural-heritage documentation and preservation, in which the entire zemi cohort participated; the per-quarter tool-introduction sequence is documented in Supplementary A3. Year 2 operated without a uniform curriculum: project-led advisory, with each student's graduation-work trajectory shaping the tool stack and supervision cadence.

The Kyoto cultural-heritage Social-Practice integration matters for the present study because it gave the cohort a shared technical-skill foundation in custom-LoRA training on student-collected fieldwork imagery (group-LoRA training, presented at the cohort's 28 October 2024 review; group composition and training parameters documented in Supplementary A3) and because three of the five longitudinal subjects' graduation projects can be traced back, in part, to capacities laid down in this practicum.

4.3 Workshop antecedent

A 2023 university-wide ComfyUI workshop at Kyoto Seika University (two sessions on 19 and 26 July 2023; 47 enrolled participants) preceded the 2024 zemi and is documented in the arXiv preprint [13] as part of the framework-genealogy claim (Section 2.4). One workshop participant—pseudonymised here as TZ, a third-year Printmaking student at Seika at the time of the workshop, who subsequently received a President's Award for his graduation work at the close of his fourth year and is now an active professional generative-AI artist working with FLUX, ComfyUI, and custom LoRAs—returned a Round 1 interview in parallel with the zemi cohort. TZ is not an embedded unit of analysis; his returns appear once in Section 6.5 as a workshop-format comparison reference for the convergent-input-integration uptake mode and otherwise inform the framework-genealogy discussion. The focal case for the multiple case study remains the 2024–2026 zemi cohort.

5 Five subjects and their trajectories

Each sketch covers: the project named in Round 1 as the one the student was most proud of from the 2024 zemi; graduation work; tool stack across the zemi–graduation arc; AI-use ratio; core mental model; trajectory category; and spectrum position. Sketches are anchored in primary-source evidence (Round 1 and Round 2 returns, Production Notes where available, supervision-moment records, practitioner reflective documentation), with verbatim passages from redacted-copy versions. A cohort-spectrum summary table is provided at the head of Section 6.

5.1 XJ — *Resonant Void* and the cumulative-deepening clean trajectory

XJ entered the 2024 zemi as one of the cohort's most technically engaged students. The seminar project she named in Round 1 as the one she was most proud of was a real-time TouchDesigner–ComfyUI cymatics pipeline; her March 2026 graduation work, *Resonant Void* (5a), expanded that pipeline into a meditation-and-stillness installation built on a co-developed StreamDiffusion + TouchDesigner stack with a custom Tibetan-mandala LoRA (100-image dataset she assembled herself) and a custom ∼80 Hz cymatics-frequency tool built for the install. The tool-stack choice emerged from a co-development conversation in which both specific-tactical and conceptual registers were active simultaneously; the LoRA training itself was XJ's independent execution. AI use is *full* in the operational sense—the work could not exist without the diffusion-model pipeline—but creative direction is XJ's.

End-of-instruction baseline (October 2024 1C). XJ's Q10 frames nodes through a cooking-recipe metaphor: "I compared nodes to a cooking recipe. Each node represents different ingredients or cooking methods, and how you combine them and their order determines the final taste and appearance of the dish (image)." The operational-fluency register the 2026 *expansive system-designer* position formalises is already present at 2024 baseline (full Q4 / Q6 verbatim in Supplementary A3).

Mental-model evolution category: *cumulative-pluralisation*—one 2024 metaphor enriched into three layered 2026 metaphors (developed in Section 6.4). Spectrum position: *expansive system-designer*—XJ's stance is the cohort's clearest articulation of AI as a co-designer and her grad work the clearest case of a tightly-coupled human–AI generative pipeline as a finished artistic system. The reflexivity flag carried in this case is temporal: substantial conceptual maturation around the Tibetan-mandala / cymatics / cosmogram framing happened *after* the supervisor's institutional transition to UM in August 2025, under attenuated-channel conditions, and is most accurately read as XJ's independent post-Seika development rather than as a co-developed artefact.

5.2 LX — *Possession / Haunting* and the capability-vs-realisation gap

LX entered the zemi with weak baseline programming and electrical-engineering knowledge (his own Round 1 self-description) and exited with *Hyōi* (*Possession / Haunting*) (5c): RPLiDAR sensor, WS2815 LED-tape vessel network from a draped-cloth torso radially outward, with a 3D-printed Noh-mask form later cut for technical and symbolic reasons. The Production Note's central conceptual formula—*kansoku* = *fukatsu* (observation = activation), elaborated as *kansoku* = *yuketsu* = *fukatsu* (observation = transfusion = activation)—is unusually philosophically explicit, scaffolded by Husserl, Barthes, Schanzer, and Noh-mask theory.

The case anchors a distinctive trajectory category, *capability-vs-realisation gap* (5c), distinct from the cumulative-deepening pattern: LX's underlying capability is high—the RPLiDAR + WS2815 integration is non-trivial cognitive engagement with tools, and the Production Note's philosophical scaffolding is operative throughout supervision, not retrospectively applied—but multiple practitioner-suggested directions toward more ambitious realisation were declined or abandoned during development. The Production Note's §F4 names a specific realisation issue candidly: floor-level visual unreadability of the LED vessels at the install, with a hard-pipe radial remedy he had no site to test. The pattern is reported as a documented observation; the question of *why* it holds is treated with interpretive restraint.

End-of-instruction baseline (October 2024 1C). LX's Q10 is hedged and tentative—"I haven't particularly thought about it; but it's similar to cooking or gardening . . . If I had to compare ComfyUI to something, that's what I'd think of." The 2024 baseline already exhibits the metaphor-resistant, programming-equivalent register the 2026 *bounded-ethical-and-pragmatic*

position formalises (full Q4 / Q6 verbatim in Supplementary A3); the ethical-authorship rationale develops on top of this 2024 substrate rather than introducing it.

Mental-model evolution category: *abandoned-metaphor-for-literal-pipeline*—hedged 2024 cooking-or-gardening metaphor → refusal of metaphor and literal-pipeline articulation at 2026 (developed in Section 6.4). His Round 2 articulates a *bounded-ethical-and-pragmatic* rationale: an ethical-authorship hesitation in which AI is anthropomorphised as a creative agent deserving credit ("I can't think I am the creator when I create with AI. It feels like I steal AI's idea") and an engineering-pragmatic limit (AI-generated code "will be Duplicated and not efficient" at multi-project scale). Post-graduation, LX is now a working programmer in Japan; offered the practitioner's career advice to deepen AI-assisted coding skills as a marketable specialisation, he chose core C++ depth instead—a choice consistent with the bounded-ethical-and-pragmatic stance the case anchors.

5.3 LY — *Sperm Whale* and the thematic-continuity-with-methodological-pivot trajectory

LY is the cohort's most artistically experienced student by entry. His seminar work included *Dystopian Feast* (AnimateDiff + Blender low-poly) and a four-minute Blender animation he named in Round 1 as the project he was most proud of—bizarre creatures with human-like traits, conceptualised through ChatGPT and Gemini brainstorming. His graduation work, *Sperm Whale* (5b), is a Unity-engine game-like piece with a hand-rigged six-character menagerie and a non-narrative-meditation register; the single-screen-with-controller format replaces the projection-in-a-dark-room format *Dystopian Feast* used.

Trajectory category: thematic-continuity-with-methodological-pivot (5b). LY himself frames the two works as "distinct directions and methodologies." The deep-ecology / spirit-world thematic line carries through; the *methodology* by which that thematic line is pursued shifts substantially. *Dystopian Feast* operates in a socio-conceptual, rational, symbolic register—in LY's own Round 2 articulation, the work of a conscious person attempting to "construct a dream through symbols to transmit a powerful signal." *Sperm Whale* operates in a non-rational, anti-symbolic register, pursuing "the tender, tactile sensation of human connection" and "the internal forms of human relationships." The medium choice (Unity rather than Blender projection) is a consequence of the methodological pivot rather than an independent medium-affordance decision: the game-engine medium serves the embodied-experience register the methodological pivot requires.

The pre-existing intellectual material LY brought to the seminar—an interest in dreams as a manifestation of the spiritual world, informed by Surrealist attitudes toward dreaming and by the contemporary internet aesthetic communities Dreamcore and Liminalcore—supplied the raw material on which his stance on AI was formed; LY's Round 2 explicitly anchors the formation of his bounded-philosophical stance temporally to the seminar itself ("This stance began to take shape during my junior year when I started participating in seminars"), establishing the seminar as the formation site for the stance rather than its first articulation. AI use is *bounded* at both the stance and realisation levels: AI was used at the brainstorming layer (a two-week phase the Round 1 returns describe explicitly) and for technical-difficulty resolution; the Unity build was hand-crafted in Blender + Unity + URP + C#; the conceptual core remained human-driven. His Round 1 L5 articulation captured the position dispositionally—"things produced by AI are akin to those emerging from a dream; yet bringing them into reality requires a body with tangible experience"—and his Round 2 consolidated it explicitly.

Production Note evidence. LY's Production Note (*Seisaku Nōto*), titled *Sperm Whale* in the header but naming the final exhibited work as *Psyconauts* in the conclusion, frames the work as a "digital space work" (*dejitaru kūkan sakuhin*) rather than a game, premised on Deep Ecol-

ogy (*shinsō seitaigaku*) and the visualisation of the *inner spiritual world of ecosystems* beyond human-centric perspectives. The Note documents an almost entirely hand-crafted realisation pipeline: hand-modelled and hand-rigged six low-poly characters in Blender, hand-implemented Unity C# dialogue and interaction logic, hand-tuned Universal Render Pipeline lighting; no generative-AI image, animation, or audio in the realised work. The Note's "lonely cells composing the ecosystem" framing is unusually philosophically explicit for an undergraduate submission, paralleling LX's Husserl / Barthes scaffolding from a different conceptual lineage. The Note documents the methodological pivot at the form-and-medium level (Blender-projection → Unity-game-engine) while preserving thematic continuity at the creative-vision level (Deep Ecology / non-human-centric / spirit-world themes carry through from *Dystopian Feast*)—supporting the Section 6.1 trajectory designation (5b: thematic-continuity-with-methodological-pivot) at the artefact-articulation layer of evidence.

End-of-instruction baseline (October 2024 1C). LY's Q10 supplies an organisational metaphor—"AI users are the boss; Nodes are company employees"—that frames AI in social-structural control-and-direction terms. The 2026 *bounded-philosophical* position relocates the same control-and-direction concern into embodied / latent-space registers through a category-different mental model (full Q4 / Q6 verbatim in Supplementary A3; Section 6.4 develops the evolutionary arc).

Mental-model evolution category: *abandoned-metaphor-for-technical-precision*—2024 boss-and-employees organisational metaphor displaced by latent-space-convergence technical articulation at 2026 (developed in Section 6.4). Spectrum position: *bounded-philosophical*. Documented practitioner interventions include Cheng's *Emissary's Guide to Worlding* [47] as vocabulary-supply (Section 6.5) and two dispositional teachings LY confirms recalling and reflecting on in his Round 2 response, both articulated by the practitioner: a *pine-trees-at-Zen-temples* example as a model for "finding the angle between the conventionally aesthetic and the merely bad" (practitioner's framing), and the dialectical formulation "anti-aesthetics is still a form of aesthetics" (also practitioner's framing). LY's Round 2 confirms both: "I certainly remember that example; I have been reflecting on it and 'chewing' over it repeatedly." The reflexivity flag in this case is biographical-affinity between supervisor and student aesthetic territory (owned in Section 3.6 / Section 7.2). Postgraduate-continuity: LY's current postgraduate research topic is "visualization of the interaction between the spiritual and material worlds", a direct lineage from *Sperm Whale* and a primary-source indicator that the bounded-philosophical stance is operative beyond the graduation transition.

5.4 LJ — *smile-detection installation* and the leap-beyond-via-project-supervision-handoff trajectory

LJ entered the zemi self-describing as having weak baseline programming familiarity and a tendency not to study actively on her own (5d). The seminar project she named in Round 1 as the one she was most proud of was conceptual rather than technical: a smile-detection installation in TouchDesigner with a recognition-as-performance frame—the work *dissolves* on detected smile, a self-dissolution-on-being-measured framing that read as one of the cohort's clearest single conceptual moves. Her graduation work extended this depth: conceptual sophistication categorically beyond what the zemi-era trajectory predicted, while technical execution operated within bounds the ambition stretched against.

Trajectory category: leap-beyond-via-project-supervision-handoff (5d). LJ's grad work was developed during the institutional-transition phase under a specific bounded-scope multi-supervisor configuration: a backup supervisor partnered with her at the project-direction layer (the conceptual / linguistic / curatorial framing of the work), while the practitioner-of-record contributed technical-architecture direction at the emotion-detection algorithm layer,

and student-self-directed template-modification handled the TouchDesigner interactive-layer implementation. The multi-supervisor arrangement is bounded to project-direction support; the broader constructionist learning across the two-year zemi remains attributable to the practitioner-of-record (Section 7 develops this distinction).

End-of-instruction baseline (October 2024 1C). LJ's Q10 is the cohort's only blank response: "Not yet…"—no metaphor articulated at the 2024 endpoint. Combined with Q4's process-attention articulation (full Q4 / Q6 verbatim in Supplementary A3), this diagnoses LJ's early-phase under-engagement and anchors the *acquired-metaphor* evolutionary category (Section 6.4): substantive mental-model articulation appears post-2024-baseline rather than evolving from a present-and-developing 2024 starting point.

Mental-model evolution category: *acquired-metaphor*—no 2024-baseline metaphor (Q10: "Not yet…") → iterative-refinement-loop framing at 2026 (developed in Section 6.4). Spectrum position: *conceptual-amplifier-only*—AI strong on conceptual / linguistic work, absent from technical execution. Her case anchors Section 6.3 mechanism (2a) with primary-source DM evidence at the 17 May 2024 TouchDesigner stuckness exchange. The framing the present paper applies is *struggle-with-good-outcome* rather than *negative case*: LJ's grad work is conceptually substantive, the project-supervision handoff produced a defensible end-state, and the case is methodologically valuable precisely because it documents a non-high-performer producing a sophisticated graduation work—providing the structured-spread evidence the constructionist outcome claim requires across the full range of student trajectories.

5.5 WS — *Shirazu Sekai* and the cumulative-deepening with realisation-quality decline trajectory

WS's seminar work culminated in an advancement-exhibition piece (*shinkyūten*, January 2025): an interactive cherry-blossom installation in Unreal Engine turning a flat screen into a window into a 3D scene. His graduation work, *Shirazu Sekai* (5e; working title earlier *Kamikakushi*), is a peripheral-vision yōkai piece premised on gaze-aversion: the yōkai become visible only when the viewer is *not* looking directly at them, mediated by a UE5 + MediaPipe gaze-detection pipeline coupled via OSC to TouchDesigner with Gaea-procedural terrain. Both works share an aesthetic through-line—screen-as-window-into-immersive-3D-scene, realistic visuals over abstract particle aesthetics—that WS's Round 2 reply explicitly named, *correcting* an earlier practitioner pre-reading that had positioned the advancement-to-graduation arc as a clean pivot away from a Refik-Anadol-inspired particle direction.

Trajectory category: cumulative-deepening with realisation-quality decline at grad-work (5e). The conceptual arc continues from advancement-exhibition piece to graduation work; the realisation under-delivers relative to the advancement-exhibition peak. The mechanism documented in WS's case (Section 6.3 mechanism 2d) is *technical-detail-fix-uptake decline* mediated by an information-disclosure-limit pattern: the practitioner offered a specific tactical fix for a camera-empty bug (yōkai visible without a viewer in frame, breaking the gaze-aversion mechanic) which was not implemented in the realised work; supervisor visibility into the development depended on what WS chose to surface, and at low-disclosure stretches the supervisor's opportunity to iterate on advice was correspondingly constrained.

End-of-instruction baseline (October 2024 1C). WS's Q10 supplies the culinary metaphor in its 2024 form: "I think of workflows as cooking steps; and various nodes as ingredients and tools." The 2024 baseline establishes both the operational-knowledge register and the culinary metaphor the 2026 articulation sustains essentially intact (full Q4 / Q6 verbatim in Supplementary A3); Section 6.4 *sustained metaphor* documents this stability.

Mental-model evolution category: *sustained metaphor*—2024 culinary register essentially unchanged at 2026, with the metaphor's role reframed from process-steps-and-materials to chef-

as-process-operator (developed in Section 6.4). Spectrum position: middle / AI-as-operations-manual—AI knows how the software operates, but creative authority stays human. His Round 2 articulates the conceptual development of *Shirazu Sekai* as "fully conceived by myself," with AI used at production stage to research yōkai related to spirit-hiding and to select gentle, friendly, non-horror creatures. The present paper takes this self-description as empirically defensible and uses *convergent-input integration* as the framing for the practitioner's 6 June 2025 lone-viewer-UX critique with peripheral-vision suggestion: the supervisor's input arrived into already-active student conceptual development at the right moment and integrated into the grad work without displacing WS's authorship of the concept (Section 6.5 develops the multi-modal register-mix uptake taxonomy).

5.6 Workshop-format comparison reference (TZ)

TZ was a third-year undergraduate in the Printmaking department at Kyoto Seika University (transitioning to fourth year) at the time of the July 2023 university-wide ComfyUI workshop documented in the arXiv preprint [13]. He is now an active professional generative-AI artist working with FLUX, ComfyUI, and custom-LoRA training. He returned a Round 1 interview in parallel with the zemi cohort; his returns are referenced once in Section 6.5 (mode 6) as a workshop-format comparison observation consistent with WS's convergent-input-integration uptake mode. TZ is not an embedded unit of analysis. His subsequent graduation work received a President's Award at the close of his fourth year at Seika—completed in the year following the workshop—and provides a documented instance in which the workshop's open-source-tool foundation met an already-developing arts-practice trajectory and contributed to its subsequent professional development.

6 Findings

6.1 The five-position spectrum at $n = 5$

The cohort produced five distinct positions on AI-as-creative-collaborator, each consistent with the student's broader creative trajectory and each defensible on its own intellectual grounds. Table 1 summarises the spectrum compactly; per-case mental-model framings, AI-use ratios, and 2024-baseline metaphors are documented at the case-sketch level in Section 5 and developed in Section 6.4.

The spectrum is not normative: no position is ranked above the others or treated as the constructionist target. The empirical claim is that the same constructionist seminar produced a structured spread of stances that range from a tightly-coupled human–AI generative pipeline (XJ) to a deliberately-bounded, philosophically-grounded constraint on AI's creative role (LY), with three distinct intermediate positions that are defensibly different from each other on identifiable grounds. This anchors the kind of differentiation the constructionist tradition predicts at the level of artefacts, extending it to articulated stance on the focal medium—available for testing at scale in subsequent work.

2024 baseline precursors. Each 2026 spectrum position has a detectable 2024-baseline precursor in the October 2024 Seika 1C Q4 (understanding-change) responses, per-case verbatim in Section 5. The five precursors reported here are restricted to the longitudinal-subject subset; the three additional 1C respondents not tracked longitudinally are not analysed for precursor mapping. The exception in structural form is LJ, whose 2024 baseline registers at the Q4 process-attention level but is blank at the metaphor level (Q10: "Not yet…")—Section 6.4 documents this as the *acquired-metaphor* evolutionary category. The spectrum did not emerge ex nihilo at 2026; it crystallised across the 18-month interval from heterogeneous 2024 articu-

Table 1: The five-position spectrum at $n = 5$.

Subj.	Spectrum position	Position in brief	Trajectory category	Graduation work
XJ	Expansive system-designer	AI as co-designer of a tightly-coupled generative pipeline	Cumulative-deepening clean (5a)	*Resonant Void* (TouchDesigner + StreamDiffusion + custom Tibetan-mandala LoRA)
WS	AI-as-operations-manual (middle)	AI knows how the software operates; creative authority stays human	Cumulative-deepening with realisation-quality decline (5e)	*Shirazu Sekai* (UE5 + MediaPipe + TouchDesigner + Gaea)
LJ	Conceptual-amplifier-only	AI strong on conceptual/linguistic work; absent from technical execution	Leap-beyond-via-project-supervision-handoff (5d)	Smile-detection installation (TouchDesigner; multi-supervisor)
LX	Bounded-ethical-and-pragmatic	AI bounded by ethical-authorship hesitation and engineering pragmatism	Capability-vs-realisation gap (5c)	*Possession / Haunting* (RPLiDAR + WS2815 + Arduino)
LY	Bounded-philosophical	AI bounded because creative authority requires embodied human experience	Thematic-continuity-with-methodological-pivot (5b)	*Sperm Whale* (Unity + C# + URP)

lations. The empirical significance for the constructionist outcome claim is developed in Section 7.2: precursors are 2024 *articulations*, not pre-zemi entry conditions, which supports a *seminar-as-formation-site* reading rather than a pre-existing-difference-amplification reading.

6.2 Two students at the bounded end with distinct rationales

The two students at the bounded end of the spectrum, LX and LY, share more than independent low-AI-use trajectories: both locate human creative authority in something AI does not fully have. They differ on what that something is. LX's rationale (Round 2) is *ethical-authorship*: AI is anthropomorphised as a quasi-creative agent ("He, maybe or She") deserving credit, and to take credit for an AI-generated artefact would be to "steal AI's idea." A second strand of his rationale is engineering-pragmatic: AI-generated code does not yet organise multi-project codebases efficiently. LY's rationale is *embodiment*, articulated dispositionally in Round 1 (AI outputs are dream-like and require "a body with tangible experience" to be brought into reality) and consolidated explicitly in Round 2: "Artificial intelligence is undeniably useful within defined parameters, yet the core of its creativity remains rooted in embodied human experience. I believe this to be a fundamental truth." LY's bounded-philosophical position is documentary-anchored at three time points: the Round 1 dream-and-body image; the Production Note's hand-crafted-with-minimal-AI realisation register for *Sperm Whale*; and the Round 2 verbatim philosophical consolidation. Different metaphysics, similar conclusion.

The pair shows that the seminar produced two bounded-end students with *thoughtful disagreement* about what AI is for, not convergent low-use trajectories driven by skill or comfort. This is among the corpus's strongest evidence that the constructionist outcome differentiation operates at the level of stance on the medium, not just at the level of project diversity.

6.3 Asymmetric-amplification: four mechanisms with three workaround paths

The cohort's longitudinal record documents that AI scaffolds non-technical-student work at the conceptual layer but realisation can fall short via four distinct mechanisms, each anchored at a specific student case. These are mechanisms by which the *amplification* of AI scaffolding

becomes asymmetric: the conceptual layer benefits while the realisation layer does not. With $n = 5$, each mechanism is anchored at a single student case and the mechanism-and-case distinction is necessarily preliminary; the typology is offered as a candidate analytic framework for the Section 7.4 UM-cohort test, not as a saturated set. The cohort also produces a workaround taxonomy alongside the failure taxonomy.

Mechanism (2a) — AI-tool-modality + sparse-training-coverage limit (LJ-anchored). Where the tool is GUI-heavy or visual-node-network with sparse training data, AI scaffolding fails. The primary-source anchor is LJ's 17 May 2024 TouchDesigner stuckness DM, where she described a GUI-node configuration she could not get working and AI advice was insufficient to bridge the gap. The positive-case end of this mechanism is XJ's 6 July 2024 all-day Docker-on-Mac co-debugging session: text-and-code modality, AI-rich training data, AI scaffolds successfully when the practitioner is also engaged in the loop. The mechanism is not "AI fails for non-technical students"; it is "AI fails when the tool's modality and training-data coverage are not aligned with what AI can scaffold."

Mechanism (2b) — student creative-direction-uptake narrowing (LX-anchored). High underlying capability + practitioner advice offered + student systematically chooses a narrower realisation than the practitioner advised. The LX case alone documents four distinct instances: an LLM-driven-spirit-mask suggestion declined; advancement-exhibition user-interaction advice not implemented; an OpenCV branching-video combinatorial-complexity workaround intent abandoned; mask-cut workaround attempts declined. The mechanism is not a failure of AI scaffolding but of practitioner-suggested ambition reaching the realisation layer. The capability-vs-realisation gap (trajectory 5c) is the per-case shape of this mechanism.

Mechanism (2c) — student medium-affordance ceiling (LY-anchored). A well-justified medium choice produces a presentation-aesthetic register cost without compromising the creative vision. *Sperm Whale* as Unity-game-engine work is the conscious decision to accept single-screen-with-controller presentation cost in service of the embodied-experience / non-narrative-meditation register the methodological pivot from *Dystopian Feast* required. The realisation-aesthetic-impact axis decouples from the creative-vision axis: the work serves its vision; the visible aesthetic-impact is lower than the projection-installation register the prior work occupied. This is not a failure mode but a documented decoupling that the spectrum analysis cannot flatten into "stronger/weaker realisation."

Mechanism (2d) — student technical-detail-fix-uptake decline (WS-anchored), mediated by information-disclosure-limit. The practitioner offers a specific tactical fix; the student declines uptake. The WS-anchored instance is the camera-empty bug at the *Shirazu Sekai* graduation work: yōkai visible without a viewer in frame, breaking the "appears when not looking" mechanic; a practitioner-offered Boolean fix was not implemented. The mediating mechanism is *information-disclosure-limit*: students who do not surface work for follow-up review do not receive iteration on advice. This is a decoupling between supervisor-input and realised-output that operates at scale wherever instructor attention is naturally rationed.

Workaround taxonomy: three documented paths. When the asymmetric-amplification limit applies, the cohort documents three workaround paths: (i) *AI-co-engineering with practitioner scaffolding* (XJ's Docker-on-Mac session); (ii) *practitioner-scaffolded high-level technical-architecture direction with research-handoff* (LJ's smile-detection emotion-detection-layer architecture, with the practitioner directing the high-level approach and the implementation detail handed off to AI / community templates); (iii) *student-self-directed template-modification when community templates exist* (LJ's TouchDesigner interactive layer, modified from a community template).

2024-baseline help-seeking dispositions as predictive of asymmetric-amplification mechanism. Each subject's October 2024 1C Q8 (help-seeking-vs-experimentation disposition)

corroborates the Section 6.3 mechanism documented at their 2026 endpoint:

- XJ — "I intend to continue experimenting ... but if I can't solve something, I'll look for various tutorials for help" → workaround positive-case (path i, AI-co-engineering eligibility).
- WS — "Experimenting to solve problems. By solving problems myself, I can learn more" → mechanism (2d) self-solving limits supervisor iteration opportunity.
- LX — "I first try to find solutions myself ... I don't like asking others for help because waiting for responses takes time" → mechanism (2b) solo-with-research declines instructor-co-engineering.
- LY — "I prefer finding alternative ways to achieve goals rather than solving problems ... Debugging when necessary isn't fun" → mechanism (2c) medium-choice avoidance of AI-execution problems.
- LJ — "Seeking help. The reason is I can quickly identify and solve unclear points" → scaffolded-articulation workaround (path ii); the project-supervision handoff worked because help-seeking is the structural condition it requires.

All five 2024 baseline dispositions are consistent with the mechanism documented at the corresponding 2026 endpoint, providing a second within-cohort triangulation parallel to the typology convergence developed in Section 6.4. The mechanism typology was constructed primarily from 2026 sources; the 2024-Q8 mapping is post-hoc consistency, not prospective prediction.

6.4 Mental-model evolution: five categorically different 2024 → 2026 arcs and the trajectory-typology cross-validation

The cohort's 2024 1C Q10 (metaphor / analogy) and Q8 (stuck — experiment vs help) entries pair with each subject's 2026 Round 1 Q8 articulation to produce a *five-category typology of mental-model evolution* across the 18-month two-timepoint pair. The five categories are not five points along a single axis; each captures a categorically different *structural form* of the 2024-to-2026 transition.

(i) Cumulative-pluralisation (XJ). One metaphor at 2024 baseline → three layered metaphors at 2026. Cooking-recipe at 2024 (Q10: "I compared nodes to a cooking recipe. Each node represents different ingredients or cooking methods, and how you combine them and their order determines the final taste and appearance of the dish (image).") is enriched into a craft-process plural at 2026: clay-being-trimmed-by-denoising, ComfyUI-as-modular-assembly-line, and image-generation-as-photographic-developing-in-a-darkroom. The 2026 metaphor set does not displace the 2024 cooking-recipe; it pluralises it. *Additive enrichment without displacement* is the structural signature.

(ii) Sustained metaphor (WS). Same culinary register at both timepoints. The 2024 Q10—"I think of workflows as cooking steps; and various nodes as ingredients and tools"—is essentially unchanged at 2026: "ComfyUI making images is like inviting a chef into your home who works with your ingredients and seasonings." The metaphor's *role* in the process shifts (process-steps + ingredients + tools → chef-as-process-operator + ingredients-as-input) but the underlying culinary register is sustained. *Stability at the conceptual / metaphorical layer* is the structural signature.

(iii) Abandoned-metaphor-for-literal-pipeline (LX). Tentative cooking-or-gardening metaphor at 2024 baseline (explicitly hedged: "I haven't particularly thought about it; but

it's similar to cooking or gardening ... If I had to compare ComfyUI to something, that's what I'd think of") → refusal of metaphor and a literal-pipeline articulation at 2026: "Load the model and LoRA, then enter the prompt in the clip field. After configuring the scheduler and setting the image size, generate the image via the sampler and decoder." The 2024 entry is already anti-metaphorical in posture; the 2026 endpoint formalises that posture into engineering pragmatism. *Posture sharpens toward narrower-but-more-controllable framings* is the structural signature.

(iv) Abandoned-metaphor-for-technical-precision (LY). Boss-and-employees organisational metaphor at 2024 baseline (Q10: "AI users are the boss; Nodes are company employees") → refusal of metaphor and a latent-space-convergence technical articulation at 2026 ("the AI is guided by the prompt within the latent space. Through multiple iterations, it filters out the noise, allowing the image to gradually converge from chaos into a clear structure. Finally, a VAE decoder is used to convert the abstract latent space data back into a pixelated image"). Both 2024 and 2026 framings are about *control and direction*, but the social-structural register pivots into a technical-precision register. *Register pivots at both work-level and mental-model-level* is the structural signature.

(v) Acquired-metaphor (LJ). No metaphor articulated at 2024 baseline (Q10: "Not yet...") → a developed iterative-refinement-loop framing at 2026 emphasising the user-side prompt → output → adjust → re-prompt loop. This is the cohort's only case of mental-model articulation *emerging* in the longitudinal window rather than being present-and-evolving from the start. *Late emergence of substantive articulation* is the structural signature.

The five categories are *structurally independent*: any single dimension (e.g., a metaphor-vs-technical-precision axis) collapses (iii) and (iv), which differ on 2024 starting register, and (v), which is about emergence rather than transition. The cohort produces a spread *across* categorically different evolutionary forms, not a spread *along* one axis.

Cross-validation with the Section 6.1 trajectory typology. The five mental-model evolutionary categories map one-to-one onto the five trajectory categories the Section 6.1 spectrum table assigns, as summarised in Table 2.

Table 2: Cross-validation of mental-model evolution categories against trajectory categories.

Subj.	Mental-model evolution	Trajectory category	Structural form (shared)
XJ	Cumulative-pluralisation	(5a) Cumulative-deepening clean	Additive enrichment without displacement
WS	Sustained metaphor	(5e) Cumulative-deepening with realisation-quality decline	Stability at the conceptual / metaphorical layer
LX	Abandoned-metaphor-for-literal-pipeline	(5c) Capability-vs-realisation gap	Posture sharpens toward narrower-but-more-controllable framings
LY	Abandoned-metaphor-for-technical-precision	(5b) Thematic-continuity-with-methodological-pivot	Register pivots at both work-level and mental-model-level
LJ	Acquired-metaphor	(5d) Leap-beyond-via-project-supervision-handoff	Late emergence of substantive articulation after supervision-handoff scaffolding

The two typologies share some 2026 source material but their primary analytic anchors are disjoint (graduation-work trajectory vs. the Q10 metaphor across the 2024–2026 pair) and their coding procedures differ. The convergence is therefore within-cohort triangulation across data sources, temporal anchors, and coding procedures by a single analyst—consistent with the Section 3.5 reflexive-thematic-analysis stance—not across-analyst inter-rater agreement, which RTA explicitly disavows. The interpretive significance for the constructionist tradition is developed in Section 7.

6.5 Multi-modal register-mix mentorship: seven uptake-modes across the cohort

The cohort's mentorship record refines an earlier analytical framing in which register-mix mentorship was characterised by per-student-uptake-mode classification—a single dominant uptake style per student. The 2026 per-student survey produced converging evidence that uptake is *multi-modal within a single student* across mentorship sources, content layers, and supervision-frequency contexts. Seven candidate uptake-modes are documented across the cohort—three of them anchored at LJ's case alone (modes 4, 5, and 7)—and the typology is offered for testing at the UM-cohort scale rather than as a saturated set:

1. **Co-decision-making** (XJ's tool-stack co-development), where supervisor and student are co-decision-makers in a single conversation that operates in both specific-tactical and conceptual registers simultaneously.

2. **Creative-direction-uptake-narrowing** (LX, four documented instances), where practitioner-suggested ambition is offered and a narrower realisation is chosen.

3. **Vocabulary-supply + register-direction + audience-engagement-reframe + dialectical-aesthetic-teaching** (LY, four intervention-types across the two-year zemi; LY's channel holds ∼47% of the cohort's DM corpus by message count). The four moments LY explicitly recalls and attributes (per Section 5.3): Cheng's *Emissary's Guide to Worlding* [47] as vocabulary-supply; audience-engagement-reframe across the zemi; the pine-trees-at-Zen-temples example as a model for "finding the angle between the conventionally aesthetic and the merely bad"; and the dialectical formulation "anti-aesthetics is still a form of aesthetics". The case-density operates at the level of *transmissible content*, not just rapport.

4. **Scaffolded-articulation** (LJ at the conceptual layer with a backup supervisor), where the conceptual / linguistic / curatorial framing of the work is co-developed across a multi-supervisor configuration.

5. **Explicit-direction-with-research-handoff** (LJ at the technical-architecture layer with the practitioner-of-record), where the high-level technical approach is directed and the implementation detail is handed off to AI or community templates.

6. **Convergent-input integration** (WS at the conceptual layer with the 6 June 2025 critique meeting active student conceptual development): practitioner input arrives into already-active student conceptual development and integrates into the work without displacing the student's authorship. The workshop-format comparison (TZ, Section 5.6) provides one consistent observation in a different pedagogical encounter; focal-case evidence is WS-only.

7. **Self-directed template-modification** (LJ at the implementation layer), which is not strictly a mentorship uptake-mode but a third strategy operating in parallel within the same student case, drawing on community-template availability rather than supervisor input.

A cross-case observation is that some single students operate multiple modes simultaneously: LJ's grad work draws on scaffolded-articulation, explicit-direction-with-research-handoff, *and* self-directed template-modification at three different content layers within the same project. This is the strongest single piece of evidence that the per-student-uptake-mode classification is under-specified by the cohort's actual mentorship record. A second cross-case observation is

register-split within a single dyad (WS): personal-rapport and advising-session registers operating in parallel within the same dyad. The pattern is present in this cohort and requires further investigation across other cohorts to determine generality.

6.6 Methodological dividend

The findings depend on triangulation across the five streams in ways worth naming. Four examples illustrate the dividend; the fourth is a *cross-validation dividend* in which two independently-constructed typologies converge.

First, the multi-modal register-mix reframe in Section 6.5. Interview returns alone would have typed LY as "attitude-only" and the others as tactical-specific. The DM record showed LY's channel actually held the highest-density mentorship in the cohort (∼47% of corpus by message count), with intense attitudinal *and* specific content-mentorship in parallel across the two-year zemi. The interview-vs-DM gap reflects what the student found *quotable* versus the *register density of supervision actually received*.

Second, the May–June 2025 synchronous critique-deployment window. Four substantive supervision letters appear in the DM corpus across a 24-day window: LY (13 May, dark-ecology endorsement); XJ (29 May, *Resonant Void* direction letter); WS (6 June, lone-viewer-UX critique); LX (6 June, one-to-one-mapping vs branching-media-tree critique). The documentary record alone reads as cohort-scheduled supervisor deployment; the per-student survey reframed the synchronicity as *turning-point-triggered*—multiple students hit graduation-work thresholds in the same window, and the practitioner's responsiveness register surfaced only when interview and survey data triangulated the documentary trace.

Third, the WS-pivot reframe documented in Section 3.4. The practitioner's pre-interview characterisation of WS's advancement-to-graduation arc as a clean pivot away from a Refik-Anadol-inspired particle direction was corrected by WS's Round 2 reply (the particle direction was an *abandoned earlier consideration*; both pieces share a 3D-window aesthetic). The Production Note Chapter 1 corroborates the cumulative-deepening reading independently. Three sources converge on a reading the practitioner-memory-only account would have got wrong; the convergence is what supports the reading.

Fourth, the Section 6.4 typology convergence and the Section 6.3 Q8 mapping are both dividends of the two-timepoint architecture: the 2024 baseline is an independent corroborative path that does not depend on the 2026 data streams from which the typologies were built.

7 Discussion

The findings support a stronger reading of the constructionist "wide walls" claim than the literature usually advances, and engage two genuine objections.

7.1 Wide walls produce stance differentiation, not just project diversity

The wide-walls claim is sometimes read as predicting only that constructionist environments support a diversity of student *projects*. The empirical demonstration in Sections 5–6 supports a stronger reading: the wide-walls space produced a structured spectrum of student *stances* on what AI is for and how it should relate to creative authority. The five positions documented at year-plus-one are not idiosyncratic per-student variations but recognisable, defensibly-different relationships with the medium, each consistent with the student's broader trajectory and articulated in interview returns and Production Notes that the analyst could not have generated on the students' behalf.

The stronger reading needs a more careful statement than "wide walls produce stance differentiation." The recalibrated framing the cohort survey supports is this: *the wide-walls space provides the conditions under which register-mix mentorship operates to amplify and articulate students' existing intellectual capacities into thoughtful, defensible, individually-different relationships with AI as a creative tool.* The pedagogy alone does not produce the outcome; the pedagogy provides the *space* in which mentorship operates, and the multi-modal register-mix mentorship documented in Section 6.5 is the active component that converts wide-walls availability into stance differentiation. This is a more cautious claim than the original wide-walls reading and more defensible at the $n = 5$ evidence base.

7.2 Engaging the prior-individual-differences objection

A genuine objection to the spectrum claim is that the differentiation reflects pre-existing individual differences amplified by the seminar rather than constructed by the pedagogy. The objection has force: students enter with different prior commitments, orientations, and dispositions; the exit spectrum could reflect those entry differences with the seminar as uniform amplifier rather than producer of differentiation.

Three considerations qualify rather than dismiss this reading. First, primary-source seminar-as-formation-site testimony: LY's Round 2 explicitly anchors the formation of his bounded-philosophical stance to the seminar itself—"This stance began to take shape during my junior year when I started participating in seminars." This is the cleanest single rebuttal in the corpus: the student names the pedagogy as the formation site for the stance, and the pre-existing intellectual material LY brought to the seminar (Surrealist attitudes toward dreaming; Dreamcore and Liminalcore aesthetics) supplied the raw material on which the stance was constructed rather than constituting the stance itself. Second, the student-side-correction pattern in Section 3.4: the cohort's interview returns substantively revise practitioner pre-readings rather than confirming them, including LY's correction of the practitioner's craft-defensive hypothesis into a substantively different philosophically-grounded position and WS's correction of the practitioner's clean-pivot framing into a cumulative-deepening reading. The students were not passive recipients of a uniform constructionist amplification of pre-existing differences; they were active articulators of stances they retrospectively attribute to the seminar; the causal direction (seminar $\rightarrow$ stance vs. prior disposition $\rightarrow$ stance) is not fully resolvable on retrospective data, and the analysis treats the seminar as the *articulation site* whether or not it is also the formation site in a stronger causal sense. Third, the calibrated framing the Section 6.5 register-mix typology supports: the mentorship documented in the DM corpus is multi-register and varies by student case rather than uniform across the cohort, and the students engaged with different combinations of those registers. The differentiation is jointly produced—pedagogy, register-mix mentorship, and amplified existing student capacity—not a pre-existing-differences artefact.

The pedagogy-and-mentorship-as-jointly-producing-the-outcome framing also has the advantage of remaining honest about the limits of the evidence. The biographical-affinity finding documented in the practitioner's per-student reflexive survey—supervisor-and-student-aesthetic-territory overlap as a non-portable enabling condition for the deepest register-mix dyad in the cohort—is named here rather than smoothed over. The Round 2 returns identified the biographical-affinity register as both dark-aesthetic commonality and a shared dialectical style of thinking—about beauty, anti-aesthetics, Zen, and Surrealism—visible in the practitioner's dispositional teachings (Section 6.5 mode 3). The resonance produced *transmissible content*—specific teachings LY can name and attribute years later—rather than just rapport. This is more useful evidence than the practitioner's pre-survey biographical-affinity hypothesis alone produced, but the contingent-feature reading still holds: biographical-affinity is a feature of the LY case rather than a generalisable design principle, and any framework drawn from this

study should flag it explicitly rather than over-rotate the constructionist outcomes claim into "supervisor's biography produced the outcome."

7.3 A second objection: is "register-mix mentorship" a renamed mentor-effect?

A second genuine objection to the findings is that the register-mix mentorship typology in Section 6.5 simply renames the well-established mentor-effect documented across arts-pedagogy and apprenticeship literatures. The objection has surface force, but the analytical move the typology supports is not the mere fact that mentorship operates in multiple registers; it is the empirical demonstration that *uptake* operates multi-modally within a single student case across mentorship sources, content layers, and supervision-frequency contexts. LJ's grad work draws on scaffolded-articulation with a backup supervisor at the project-direction layer, explicit-direction-with-research-handoff with the practitioner-of-record at the technical-architecture layer, and self-directed template-modification at the implementation layer—three uptake-modes operating in parallel within the same project. This is a finer-grained characterisation of pedagogical uptake than a mentor-effect framing supports, and it makes the multi-supervisor finding (Section 5.4) tractable as a bounded scaffolding feature: the project-direction support and the constructionist learning are at different content layers, and the cohort's per-case uptake-mode distribution is empirically reconstructable from the five data streams in ways a unitary mentor-effect framing cannot accommodate.

7.4 Future research

The principal future-research direction internal to the present project is a parallel cohort of postgraduate students at the University of Macau who received a closely-aligned protocol in April–May 2026, and whose returns will provide a substantially larger comparison sample at a different institution, in a different cultural setting, with a different tool ecosystem, and under the same constructionist pedagogical principles. The UM protocol preserves the Seika instrument's 12 core questions and adds three items designed against the Seika findings: a convergent-input-integration probe (targeting the Section 6.5 uptake-mode 6), an information-disclosure-pattern probe (targeting Section 6.3 mechanism 2d), and an asymmetric-amplification-realisation-limit probe (targeting the Section 6.3 four-mechanism typology). The UM data will test whether the multi-modal uptake-mode typology in Section 6.5, the four-mechanism asymmetric-amplification finding in Section 6.3, and the mental-model spread in Section 6.4 transfer at scale. Three pre-registerable predictions follow from the Seika findings: the UM cohort should produce a recognisable spectrum of stances at exit-from-course; the four asymmetric-amplification mechanisms should appear in some distribution across the larger cohort; and the multi-modal register-mix uptake typology should hold under conditions where individual instructor attention is naturally rationed across the larger cohort.

8 Limitations

Seven limitations warrant explicit naming.

Single-site, single-instructor, $n = 5$ longitudinal subjects. The empirical site is one undergraduate zemi at one Japanese art university taught by one practitioner. Generalisation beyond the case is constrained accordingly: findings extend to the theoretical propositions of constructionist learning theory and practitioner inquiry [39], not to a population of seminars or students. Cross-cohort generalisation of the spectrum claim itself requires the larger UM cohort sample analysed in subsequent work (Section 7.4).

Single-researcher RTA. The analysis is conducted by a single researcher who is also the practitioner of record. The reflexive thematic analysis framework [41] explicitly accommodates this stance (Section 3.6); the limitation is named for readers operating in a coding-reliability paradigm rather than as a concession that RTA delivers a weaker analytical claim.

Self-report bias in interview returns. The Round 1 and Round 2 interviews are self-report data subject to social-desirability bias and imperfect recall. The instructor-researcher power dynamic is partially mitigated by interview administration after course completion, by explicit disclosure that participation has no bearing on past or future academic standing, and by the student-side-correction pattern in Section 3.4, which provides empirical evidence that participants did not write diplomatically. The bias is not eliminated; it is acknowledged.

Cross-linguistic translation. Interviews were administered in English to participants for whom English is a second or third language; passages in Japanese and Chinese were permitted where participants felt they could express ideas more precisely. Translation into English for the analysis introduces inherent risk of nuance, metaphor, and culturally specific meaning being lost or distorted. The Production Notes are bilingual (Japanese with English translation); supervision-record verbatim quotation uses redacted-copy versions of the original language with English gloss. The limitation applies particularly to the LX bounded-ethical-and-pragmatic framing, which surfaces an anthropomorphisation of AI in English ("He, maybe or She") whose cross-linguistic resonance the analysis cannot fully assess.

Retrospective-articulation timing. The Round 1 and Round 2 interviews were conducted approximately eighteen months after the end of Semester 1 of Year 1 (October 2024 → April–May 2026) and one month after the cohort's graduation exhibition. The findings about each subject's 2026 stance are therefore about *retrospective student articulation* of stances on AI-as-creative-collaborator, not about real-time stance-formation during the seminar. The two-timepoint longitudinal architecture described in Section 3.3—pairing the 2024 1C end-of-instruction qualitative survey against the 2026 interview returns—provides the closest available substitute for contemporaneous interviewing across the arc, but it does not constitute continuous in-arc data collection at every intermediate phase, and the inferential claim available is therefore *2024-baseline → 2026-endpoint evolution* rather than *real-time-stance-formation-during-the-seminar*. The DM corpus and Production Notes provide additional intermediate anchors at supervisor-engagement and graduation-work-formulation phases respectively, but they do not deliver standardised cohort-wide qualitative-survey instruments between the 2024 baseline and the 2026 endpoint.

Multi-supervisor handoff (LJ). The leap-beyond-via-project-supervision-handoff trajectory (Section 5.4) is bounded to LJ's case. The backup-supervisor partnership operated at the project-direction layer only; the constructionist learning across the seminar arc remains attributable to the practitioner-of-record. This is a per-case scaffolding feature, not a generalisable design principle.

Biographical-affinity (LY). The supervisor–student aesthetic-territory overlap discussed in Section 7.2 enabled a deeper register-mix dyad than other cohort dyads. This is a contingent enabling condition, not a generalisable design principle.

9 Conclusion

This paper documents a five-position spectrum of AI-as-creative-collaborator stances populated by five students from a single constructionist two-year zemi at Kyoto Seika University. The same zemi, run on the same constructionist pedagogical principles around the same Semester-1 ComfyUI anchor (extended subsequently to a mixed open-source-and-proprietary ecosystem), produced five recognisably-different stances (Table 1). The differentiation is structured rather than

idiosyncratic and persists into post-graduation creative practice. The spectrum is *triangulated within-cohort* by an independently-constructed mental-model evolution typology over the 2024 1C → 2026 Round 1 longitudinal pair: the two typologies converge one-to-one onto the same five-position structure. A four-mechanism asymmetric-amplification finding documents how AI scaffolding can fail at the realisation layer despite succeeding at the conceptual layer, with each mechanism corroborated by its subject's 2024-baseline help-seeking disposition; a three-path workaround taxonomy documents what scaffolding can recover. A multi-modal register-mix mentorship typology documents seven uptake-modes operating simultaneously across mentorship sources, content layers, and supervision-frequency contexts.

What the framework offers other AI-art-pedagogy researchers and practitioners is a calibrated reading of the constructionist "wide walls" claim: the wide-walls space provides the conditions, and multi-modal register-mix mentorship is the active component that converts that space into stance differentiation. The methodological argument made alongside the empirical findings—that pairing an end-of-instruction qualitative-survey baseline with a longitudinal follow-up interview round, triangulated against DM records, Production Notes, and practitioner interpretive content, produces analytic findings (including the cross-validation dividend) none of the sources alone supports—is offered as a transferable resource for practitioner-inquiry case-study research in AI in higher education.

Two future-research directions warrant naming. First, a parallel cohort of postgraduate students at the University of Macau received a closely-aligned protocol in April–May 2026, in a different institutional context, cultural setting, and tool ecosystem. The UM data will test whether the uptake-mode typology, the four-mechanism asymmetric-amplification finding, and the mental-model spread transfer at scale. Second, the durability of the spectrum into subsequent academic and professional stages: one subject's postgraduate research topic—"visualization of the interaction between the spiritual and material worlds"—continues the bounded-philosophical lineage, a primary-source datum that the spectrum survives the graduation transition; the durability question across other cohort members remains open and is addressed in the present study's continuing longitudinal follow-up.

Declarations

Ethics approval and consent to participate

The study was approved by the Kyoto Seika University Ethics Review Committee. The research was conducted in accordance with the Kyoto Seika University Research Ethics Regulations and the Kyoto Seika University Guidelines for Research Involving Human Subjects. The research activities reported here were supported by the President's Designated Research Grant (*Gakuchō-Shitei-Kenkyū*, 2024). All participants provided informed written consent for interview content, Production Note excerpting, and verbatim quotation of supervision records.

Consent for publication

All participants provided consent for anonymised reproduction of interview content and Production Note excerpts; supervision-record verbatim quotation uses redacted-copy versions of original artefacts.

Data availability

The datasets generated during and/or analysed during the current study are not publicly available due to confidentiality commitments made to participants and the presence of identifying

information in the interview transcripts, Discord supervision records, and Production Notes, but are available from the corresponding author on reasonable request.

Competing interests

The author declares no competing interests.

Funding

University of Macau Start-Up Research Grant SRG2025-00057-FAH (2025–2027). Earlier phases of the work were supported by Kyoto Seika University's President's Designated Research Grant (*Gakuchō-Shitei-Kenkyū*, 2024).

Author contributions

Single author. The author conceived and designed the study, conducted data collection and analysis, and drafted the manuscript.

Use of large language models

Large language models (Claude, ChatGPT) were used for translation support of Japanese and Chinese interview passages, for drafting assistance, and for editorial copy-editing. LLMs did not function as independent coders or as adjudicators of theme structure or category assignment; all analytical decisions about coding, theme construction, and case-study interpretation were made by the author. The author retains full responsibility for the manuscript's analytical claims and final text.

References

[1] Mitchel Resnick. *Lifelong Kindergarten: Cultivating Creativity Through Projects, Passion, Peers, and Play*. MIT Press, Cambridge, MA, 2017.

[2] Seymour Papert. *Mindstorms: Children, Computers, and Powerful Ideas*. Basic Books, New York, 1980.

[3] Seymour Papert. Situating constructionism. In Idit Harel and Seymour Papert, editors, *Constructionism*, pages 1–11. Ablex Publishing, Norwood, NJ, 1991.

[4] Yasmin B. Kafai. Playing and making games for learning: Instructionist and constructionist perspectives for game studies. *Games and Culture*, 1(1):36–40, 2006. doi: 10.1177/1555412005281767.

[5] James Hutson and Pace Cotroneo. Generative AI tools in art education: Exploring prompt engineering and iterative processes for enhanced creativity. *Metaverse*, 4(1):1–14, 2023. doi: 10.54517/m.v4i1.2164.

[6] John V. Pavlik and Olivia M. Pavlik. Art education and generative AI: An exploratory study in constructivist learning and visualization automation for the classroom. *Creative Education*, 15(4):601–616, 2024. doi: 10.4236/ce.2024.154037.

[7] Kıvanç Tatar, Petter Ericson, Kelsey Cotton, Paola Torres Núñez del Prado, Roser Batlle-Roca, Beatriz Cabrero-Daniel, Sara Ljungblad, Georgios Diapoulis, and Jabbar Hussain. A shift in artistic practices through artificial intelligence. *Leonardo*, 57(3):293–297, 2024. doi: 10.1162/leon_a_02523.

[8] Sonia Sáez-Velasco, Mario Alaguero-Rodríguez, Vanesa Delgado-Benito, and Sonia Rodríguez-Cano. Analysing the impact of generative AI in arts education: A cross-disciplinary perspective of educators and students in higher education. *Informatics*, 11 (2):37, 2024. doi: 10.3390/informatics11020037.

[9] Younjung Hwang and Yi Wu. Graphic design education in the era of text-to-image generation: Transitioning to contents creator. *International Journal of Art & Design Education*, 44(1):239–253, 2025. doi: 10.1111/jade.12558.

[10] Yihan Jiang, Yujiao Fan, and Zifeng Liu. Generative AI in art education: A systematic review of research trends, tool applications, and outcomes (2019–2025). *Education Sciences*, 16(1):47, 2026. doi: 10.3390/educsci16010047.

[11] Zhishu Fang. Integrating generative AI in higher art education: A systematic review of tools, pedagogies, and practices. *SN Computer Science*, 7(2):215, 2026. doi: 10.1007/s42979-026-04788-x.

[12] Saebom Kim and Jeremy Breaden. The zemi in Japanese higher education: Continuity and contemporary significance. *Asia Pacific Journal of Education*, 2026. doi: 10.1080/02188791.2026.2632165. Advance online publication.

[13] Atticus Sims. From creation to curriculum: Examining the role of generative AI in arts universities. arXiv:2412.16531, 2024. Working paper.

[14] Seymour Papert. Hard fun, 2002. URL `http://www.papert.org/articles/HardFun.html`. Bangor Daily News.

[15] Yasmin B. Kafai and Quinn Burke. *Connected Code: Why Children Need to Learn Programming*. MIT Press, Cambridge, MA, 2014.

[16] Ken Kahn and Niall Winters. Constructionism and AI: A history and possible futures. *British Journal of Educational Technology*, 52(3):1130–1142, 2021. doi: 10.1111/bjet.13088.

[17] Ilya Levin, Alexey Semenov, and Mikael Gorsky. Smart learning in the 21st century: Advancing constructionism across three digital epochs. *Education Sciences*, 15(1):45, 2025. doi: 10.3390/educsci15010045.

[18] Zhixiu Lin, Yun Dai, and Oi-Lam Ng. Constructionism in K-12 AI literacy education: A systematic review of pedagogical designs, student outcomes, and learning mechanisms. *Journal of Educational Computing Research*, 63(7–8):1748–1781, 2025. doi: 10.1177/07356331251360442.

[19] Donald A. Schön. *The Reflective Practitioner: How Professionals Think in Action*. Basic Books, New York, 1983.

[20] Donald A. Schön. *Educating the Reflective Practitioner: Toward a New Design for Teaching and Learning in the Professions*. Jossey-Bass, San Francisco, 1987.

[21] Marilyn Cochran-Smith and Susan L. Lytle. Relationships of knowledge and practice: Teacher learning in communities. *Review of Research in Education*, 24:249–305, 1999. doi: 10.3102/0091732X024001249.

[22] Marilyn Cochran-Smith and Susan L. Lytle. *Inquiry as Stance: Practitioner Research for the Next Generation*. Teachers College Press, New York, 2009.

[23] Michael Polanyi. *Personal Knowledge: Towards a Post-Critical Philosophy*. University of Chicago Press, Chicago, 1958.

[24] Michael Polanyi. *The Tacit Dimension*. Doubleday, Garden City, NY, 1966.

[25] Graeme Sullivan. *Art Practice as Research: Inquiry in the Visual Arts*. Sage Publications, Thousand Oaks, CA, 2005.

[26] Linda Candy. Practice based research: A guide. Technical Report CCS Report No. 1.0,

Creativity and Cognition Studios, University of Technology Sydney, 2006.

[27] Robin Nelson. *Practice as Research in the Arts (and Beyond): Principles, Processes, Contexts, Achievements*. Palgrave Macmillan, Cham, Switzerland, 2nd edition, 2022. doi: 10.1007/978-3-030-90542-2.

[28] James Hutson and Pace Cotroneo. Praxis and augmented creativity: A case study in the use of generative artificial intelligence (AI) art in the digital art classroom. *The International Journal of Technologies in Learning*, 31(1):113–132, 2023. doi: 10.18848/2327-0144/CGP/v31i01/113-132.

[29] Jennie Black and Tom Chaput. A discussion of artificial intelligence in visual art education. *Journal of Computer and Communications*, 12(5):71–85, 2024. doi: 10.4236/jcc.2024.125005.

[30] Henriikka Vartiainen, Matti Tedre, and Ilkka Jormanainen. Co-creating digital art with generative AI in K-9 education: Socio-material insights. *International Journal of Education Through Art*, 19(3):405–423, 2023. doi: 10.1386/eta_00143_1.

[31] Henriikka Vartiainen and Matti Tedre. Using artificial intelligence in craft education: Crafting with text-to-image generative models. *Digital Creativity*, 34(1):1–21, 2023. doi: 10.1080/14626268.2023.2174557.

[32] Hexuan Zhou, Jichen Zhu, Michael Mateas, and Noah Wardrip-Fruin. The eyes, the hands and the brain: What can text-to-image models offer for game design and visual creativity? In *Proceedings of the 19th International Conference on the Foundations of Digital Games (FDG '24)*, pages 1–13. ACM, 2024. doi: 10.1145/3649921.3650001. Article 23.

[33] Anil R. Doshi and Oliver P. Hauser. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28):eadn5290, 2024. doi: 10.1126/sciadv.adn5290.

[34] Evan F. Risko and Sam J. Gilbert. Cognitive offloading. *Trends in Cognitive Sciences*, 20(9):676–688, 2016. doi: 10.1016/j.tics.2016.07.002.

[35] Kevin Crowston and Francesco Bolici. Deskilling and upskilling with AI systems. *Information Research*, 30(iConf):1009–1023, 2025. doi: 10.47989/ir30iConf47143.

[36] Jing Chen, Nur Azlina Mohamed Mokmin, and HanJun Su. Integrating generative artificial intelligence into design and art course: Effects on student achievement, motivation, and self-efficacy. *Innovations in Education and Teaching International*, 62(5):1431–1446, 2025. doi: 10.1080/14703297.2025.2503857.

[37] Mustafa Sarıtepeci and Hatice Yıldız Durak. Effectiveness of artificial intelligence integration in design-based learning on design thinking mindset, creative and reflective thinking skills: An experimental study. *Education and Information Technologies*, 29(18):25175–25209, 2024. doi: 10.1007/s10639-024-12829-2.

[38] Jiawei Wang, Jingru Zhang, and Yuan Kang. Evaluating the effects of personalised learning on AI-assisted design performance, creative self-efficacy, and engagement of high- and low-performing college design students. *Journal of Computer Assisted Learning*, 42(2):e70221, 2026. doi: 10.1002/jcal.70221.

[39] Robert K. Yin. *Case Study Research and Applications: Design and Methods*. Sage, Thousand Oaks, CA, 6th edition, 2018.

[40] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006. doi: 10.1191/1478088706qp063oa.

[41] Virginia Braun and Victoria Clarke. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, 11(4):589–597, 2019. doi: 10.1080/2159676X.2019.1628806.

[42] Virginia Braun and Victoria Clarke. *Thematic Analysis: A Practical Guide*. Sage, London, 2022.

[43] Eric Dahlin. Email interviews: A guide to research design and implementation. *International Journal of Qualitative Methods*, 20:16094069211025453, 2021. doi: 10.1177/16094069211025453.

[44] Lokman I. Meho. E-mail interviewing in qualitative research: A methodological discussion. *Journal of the American Society for Information Science and Technology*, 57(10):1284–1295, 2006. doi: 10.1002/asi.20416.

[45] Michelle Amri, Christina Angelakis, and Dilani Logan. Utilizing asynchronous email interviews for health research: Overview of benefits and drawbacks. *BMC Research Notes*, 14 (1):148, 2021. doi: 10.1186/s13104-021-05547-2.

[46] Rob Hawkes and John Adamson. Challenges and opportunities for expatriate faculty teaching graduation seminar classes in Japanese universities. *Asian-Pacific Journal of Second and Foreign Language Education*, 7:19, 2022. doi: 10.1186/s40862-022-00176-z.

[47] Ian Cheng. *Emissary's Guide to Worlding*. Metis Suns, New York, 2018. Edited by V. So and N. Sung.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryLaTeXSourceR2.zip