Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 2721–2727
https://doi.org/10.54941/ahfe1008496

Beyond Tools: A Task-Based Human-AI Studio Scaffold for Critical Thinking in Architectural Education

Silvia Albano

Xi'an Jiaotong - Liverpool University, Suzhou, Jiangsu 215123, China

ABSTRACT

Generative AI is increasingly present in architectural education, yet studio use often remains tool-centric, privileging visual acceleration over design reasoning. This paper reframes AI integration as a Human-AI interaction and UX-for-learning problem: educational value depends less on access to image generators than on interaction structures that require students to compare alternatives, justify selections, and reflect on consequences. Drawing on Human-AI interaction guidelines, creativity support tools research, architectural pedagogy, and design-cognition literature, the paper proposes a six-exercise studio scaffold covering conceptual exploration, iterative development, material and environmental simulation, diagram-to-atmosphere visualization, final visual portfolio enhancement, and AI-generated video presentation. Each exercise embeds lightweight process evidence through comparative outputs, captions, storyboards, and short written reflections. The paper is theoretical and design-oriented, but it also discusses some preliminary anonymized student exemplars as illustrative cases of framework applicability rather than empirical proof of learning gains. The main contribution is a repeatable scaffold that turns AI use into visible decision-making episodes and links ordinary studio outputs to critical-thinking dimensions such as interpretation, analysis, evaluation, inference, explanation, and self-regulation. The paper concludes by outlining a future empirical design for controlled comparison within parallel studio groups.

Keywords: Generative AI, Architectural education, Critical thinking, Human-AI interaction, AI-design studio workflow

INTRODUCTION

Generative AI has rapidly entered architectural education, particularly in concept generation, image synthesis, and representational experimentation. In many studio settings, however, adoption remains strongly tool-centric: students use AI to accelerate visualization, polish atmospheres, or produce persuasive images, while the underlying design reasoning remains weakly articulated. This creates a pedagogical risk. If studio learning is reduced to increasingly plausible outputs, students may converge too early on visually convincing solutions, accept machine suggestions without sufficient verification, and confuse representational fluency with architectural judgement.

Received October 9, 2026; Revised November 5, 2026; Accepted November 17, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This risk is especially significant in the architectural studio, where learning has traditionally depended on iterative judgement, reflective action, and the comparison of alternatives. Design is not only a matter of producing form; it is also a process of framing problems, testing propositions, and revising decisions in response to emerging constraints and meanings. From this perspective, the core issue is not whether AI can generate images, but whether studio tasks can be structured so that intent, selection, critique, and revision remain visible and assessable while AI is in use.

This paper addresses that issue by proposing a task-based Human-AI studio scaffold derived from six exercises already designed for student use. The scaffold treats AI not as a shortcut to final imagery but as a structured provocateur within a learning workflow. The paper asks two questions: which task-level interaction requirements best support critical thinking when students use generative AI in studio, and which traces can support assessment of reasoning without reducing studio work to image quality alone? The contribution is theoretical rather than empirically validated, but it is strengthened by reference to some preliminary anonymized student exemplars that illustrate how different degrees of reasoning visibility emerge in practice.

LITERATURE REVIEW

The argument is grounded in four complementary bodies of literature. First, Human-AI interaction research emphasizes transparency, user control, feedback, and recovery from error as central conditions for appropriate reliance on AI systems (Amershi et al., 2019). This is closely related to the classic distinction between use, misuse, disuse, and abuse in automation studies (Parasuraman and Riley, 1997). In educational contexts, these ideas suggest that students should not encounter AI as an unquestioned answer engine, but as a system whose outputs must be inspected, compared, and challenged.

Second, creativity support tools research argues that digital systems are most valuable when they widen exploration and strengthen human choice rather than simply accelerate production (Shneiderman, 2007). This aligns with architectural pedagogy. Cross (1982) describes design as a distinctive way of knowing grounded in synthesis and judgement, while Schon (1983) frames practice as a reflective conversation with a situation. Suwa and Tversky (1997) similarly show that sketches are not passive records but active cognitive devices through which reinterpretation and discovery occur. Taken together, these perspectives imply that AI is educationally meaningful only when it expands reflective dialogue instead of compressing it.

Third, architecture-specific literature has already identified both opportunities and tensions in AI-supported design. Castro Pena et al. (2021) show that AI is especially active in conceptual design, where it can support alternative generation and early exploration. More recent studio-based work suggests that AI can enhance learning when used in parallel design groups, but also that comparison and reflective selection remain essential if the process is not to collapse into surface-level image production (Agkathidis et al., 2024).

Fourth, broader higher-education literature has warned that generative AI may widen access to ideas while simultaneously weakening judgement if

assessment frameworks are not redesigned around verification, explanation, and responsible use (Yan et al., 2024). In architecture, where design quality is open-ended and visually persuasive output can mask conceptual weakness, the need for such scaffolds is particularly acute. What remains insufficiently developed is a task-based structure that embeds critical reflection within ordinary studio work rather than adding it afterwards as external commentary.

This perspective also helps reposition the debate away from the simplistic opposition between AI as mere tool and AI as autonomous partner. In studio teaching, the more productive question is whether AI can be configured as a pedagogical interlocutor that expands possibility while leaving judgement accountable to the student. That position is consistent with the reflective traditions of design education, which do not eliminate external aids, but require that such aids remain open to reinterpretation, critique, and deliberate revision.

METHODOLOGY

This study adopts a theoretical, design-oriented methodology based on conceptual synthesis and interaction specification. Its aim is not to demonstrate causal improvement in student performance, but to define a pedagogical framework that can later be tested empirically. The primary analytical object is the individual exercise, understood as a bounded interaction episode containing a design intention, an AI-mediated production step, a required act of human judgement, and a residual trace that can be reviewed for assessment.

The six exercises are read as a sequence of alternating divergent and convergent moments: conceptual exploration, iterative development, material and environmental simulation, diagram-to-atmosphere visualization, final visual portfolio enhancement, and AI-generated video presentation. Across these tasks, the analysis focuses on whether students are required to compare options, select directions, justify progression, and communicate design intent in ways that make reasoning legible. The mapping of these traces to critical-thinking dimensions draws on Facione's (1990) framework of interpretation, analysis, evaluation, inference, explanation, and self-regulation.

In addition to the conceptual analysis, some preliminary student submissions produced through the same workflow are discussed as illustrative exemplars. These are not treated as a formal dataset and are not used to support claims of validated educational effect since the testing happened during the definition and finetune of the workflow. Instead, they serve as anonymized cases that test the practical legibility of the scaffold and reveal the uneven depth with which reflection may appear in student work.

Methodologically, the paper therefore combines framework design with small-scale illustrative reading. The conceptual contribution lies in specifying the interaction logic of the scaffold; the illustrative contribution lies in showing how that logic appears, or fails to appear clearly, in student outputs. This distinction is important because it allows the paper to remain academically cautious while still drawing value from real studio material. The student figures are included as evidence of pedagogical legibility, not as proof of causal impact.

FINDINGS

The first finding is that the six exercises can be reformulated as a repeatable sequence of decision-making episodes.

Table 1: Six exercises.

Exercise	Student Activity	Description	Cognitive Skill
01	Conceptual Exploration	Multiple AI-generated readings; alternatives compared; one direction chosen.	Interpretation · Inference
02	Iterative Development	Prompts adjusted deliberately across iterations; each revision judged as progress or detour.	Analysis · Evaluation
03	Material & Environmental Simulation	Material and environmental options tested; choices tied to site, atmosphere, construction.	Explanation
04	Diagram-to-Atmosphere	Plans and massing translated into atmospheric perspectives.	Translation
05	Portfolio Enhancement	Final visuals paired with captions; portfolio reads as an argued proposal.	Evaluation · Explanation
06	Video Presentation	Storyboard and synopsis become a short AI-generated video; project synthesized.	Synthesis · Self-regulation

In Exercise 1, students generate multiple conceptual interpretations and select one for further development; in Exercise 2, they refine the chosen direction through prompt-adjusted iterations; in Exercise 3, they test material and environmental alternatives; in Exercise 4, they translate plans or massing into atmospheric perspectives; in Exercise 5, they pair final visuals with captions; and in Exercise 6, they synthesize the project through a short video, storyboard, and synopsis. Across all six stages, AI is most educationally productive when it is embedded in the recurring rhythm of comparison, selection, justification, and revision. This sequence matters because it shifts the educational focus from output quality to judgement quality. Comparative generation widens the search space; explicit selection preserves human agency; brief reflection makes reasoning visible without overburdening students; and captions or storyboards turn final representations into arguments rather than decorative endpoints. In this sense, the scaffold operationalizes AI-enhanced critical thinking not as a vague aspiration but as a set of observable task behaviours aligned to interpretation, evaluation, explanation, and self-regulation. The preliminary student exemplars support this argument in a limited but useful way. The strongest exemplar showed a clear progression from alternative concept directions to a deliberately chosen memorial strategy, followed by iterative refinement and explicit material reasoning linked to geology, temporality, and narrative. A second exemplar showed a well-structured visual chain from concept to development, materials, atmosphere, and final proposal, but with lighter reflection and weaker explanation of why certain options

were preferred. A third exemplar followed the same sequence more generically, demonstrating that workflow compliance alone does not guarantee analytical depth. Figures 1 and 2 illustrate two different levels of reasoning visibility within this framework. Figure 1 presents an anonymized student board in which AI use is organized as a staged process from conceptual iterations to atmosphere and final proposal. The board makes progression legible and therefore supports review, yet the critical rationale behind transitions remains relatively implicit. Figure 2 shows a stronger case in which AI-supported concept development is reintegrated into architectural drawings, sections, and elevations. Here the workflow moves beyond image generation toward a more explicit architectural translation of intent.

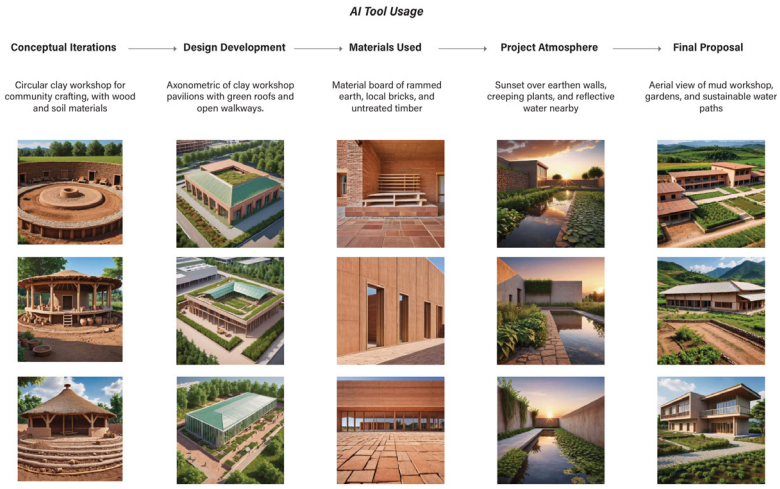


Figure 1: Anonymized student exemplar showing a staged AI-supported workflow from conceptual iterations to materials, atmosphere, and final proposal. The sequence makes progression visible, but the evaluative rationale remains comparatively implicit.

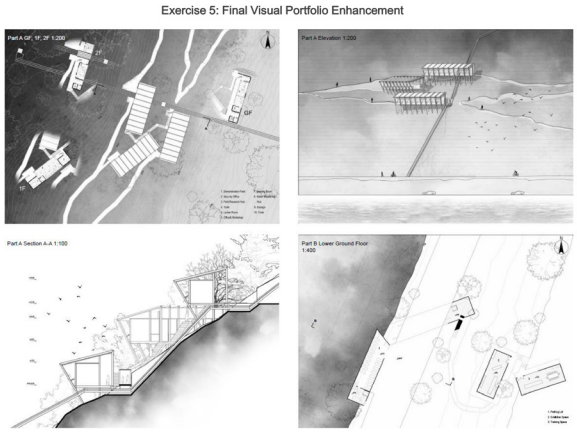


Figure 2: Anonymized student exemplar in which AI-supported concept development is translated into architectural drawings, sections, elevations, and site organization. Here the workflow moves beyond image generation toward explicit architectural articulation.

These differences are important. They suggest that the value of the scaffold lies not simply in asking students to use AI across several steps, but in requiring forms of reflection that prevent the workflow from becoming a record of visual progression only. Where comparison and explanation are explicit, tutors can reconstruct a chain of judgement from ordinary studio artefacts. Where reflection remains minimal, the same artefacts document activity without fully evidencing critical thinking. The pedagogical challenge, therefore, is not merely to insert AI into studio, but to calibrate the interaction requirements that keep architectural judgement visible. A second finding follows for assessment. Because the workflow already asks for comparative sets, concise reflections, captions, and storyboards, it produces evidence traces that can be reviewed without building an entirely separate assessment apparatus. This is significant for studio pedagogy, where excessive documentation can become counterproductive. The scaffold therefore offers a practical middle position: it does not deny the importance of visual quality, but it relocates assessment toward the reasoning that shapes those visuals.

At a finer level, the six exercises distribute critical-thinking demands across the studio timeline. Conceptual exploration emphasizes interpretation and inference because students must read differences between alternatives and anticipate which direction is worth developing. Iterative development foregrounds analysis and evaluation because students alter prompts intentionally and judge whether revisions constitute progress. Material and environmental simulation extends explanation by forcing students to connect visual choices to site, atmosphere, and constructional meaning. Diagram-to-atmosphere visualization requires translation between abstract and experiential registers. Portfolio enhancement and video presentation then ask students to synthesize and communicate the project as a coherent argument. The scaffold is thus cumulative: later outputs are stronger when earlier choices have been explicitly reasoned.

CONCLUSION AND FUTURE WORK

This paper has argued that the educational question surrounding generative AI in architecture is not simply whether the technology is available, but how interaction with it is staged within studio pedagogy. When AI is used only to accelerate representation, it risks encouraging premature convergence and weakly examined choices. When embedded in a task-based scaffold, it can support a more explicit design process in which students compare alternatives, justify decisions, and communicate intent through ordinary studio outputs.

The framework remains theoretical and should not be treated as empirically validated. The anonymized student exemplars included here demonstrate applicability rather than effectiveness: they show that the scaffold can structure student work and reveal different degrees of reasoning visibility, but they do not prove improved critical thinking outcomes. Their main value is diagnostic. They clarify where reflective prompts are working and where they still need strengthening.

Future work should test the scaffold in parallel studio groups working on the same brief under a shared rubric, combining tutor review with coded analysis of process traces. Longitudinal application would also help determine

whether habits of comparison, explanation, and self-regulation transfer into later studio tasks without explicit prompting. More broadly, the framework suggests that generative AI should be positioned in architectural education not as a replacement for design thinking, but as a structured interlocutor through which design thinking can be made more explicit, contestable, and teachable.

REFERENCES

Agkathidis, A., Song, J. and Symeonidou, C. (2024). AI-assisted design and learning in parallel design studio groups. In Proceedings of eCAADe 2024, Vol. 2, pp. 619–628.

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P.N., Inkpen, K., Teevan, J., Kikin-Gil, R. and Horvitz, E. (2019). Guidelines for human-AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, pp. 1–13.

Castro Pena, M.L., Pena, M. and others (2021). Artificial intelligence applied to conceptual design: A review of its use in architecture. Automation in Construction, 124, 103550.

Cross, N. (1982). Designerly ways of knowing. Design Studies, 3(4), pp. 221–227.

Facione, P.A. (1990). Critical thinking: A statement of expert consensus for purposes of educational assessment and instruction. Newark, CA: The Delphi Report.

Parasuraman, R. and Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), pp. 230–253.

Schon, D.A. (1983). The Reflective Practitioner: How Professionals Think in Action. New York: Basic Books.

Shneiderman, B. (2007). Creativity support tools: Accelerating discovery and innovation. Communications of the ACM, 50(12), pp. 20–32.

Suwa, M. and Tversky, B. (1997). What do architects and students perceive in their design sketches? A protocol analysis. Design Studies, 18(4), pp. 385–403.

Yan, L., Greiff, S. and Gasevic, D. (2024). Promises and challenges of generative artificial intelligence for human learning. Nature Human Behaviour, 8, pp. 1839–1850.