The Impact of Generative AI on Architectural Design Education: Creativity, Learning Outcomes, Critical Thinking, and Academic Integrity in the Studio

Ar. Snehal Kotadiya[1], Ar. Alpesh Golakiya[2]
[1,2] *P. P. Savani University*

***Abstract*—Generative artificial intelligence (GenAI) — including text-to-image diffusion models, large language models, and bespoke architectural plug-ins — has moved from novelty to near-standard presence in architectural design studios within a few short years. This paper synthesizes recent empirical and conceptual literature to examine four interlocking questions: how GenAI affects student creativity, how it shapes measurable learning outcomes, what it does to critical thinking and independent design judgment, and what new academic integrity challenges it introduces to a discipline that has always depended on visual authorship. The review finds that GenAI tools are broadly experienced by students as creativity-supportive during early ideation, associated with measurable gains in design form and aesthetic quality, but carry a documented risk of eroding critical thinking and design agency when used without deliberate pedagogical scaffolding, particularly in later design stages where originality and judgment matter most. It also finds that architecture's existing culture of iteration, referencing, and collective critique leaves it especially exposed to unresolved questions of authorship and attribution. The paper concludes by outlining a stage-differentiated, transparency-based framework for integrating GenAI into architectural curricula in ways that preserve rather than displace the core competencies design education exists to build.**

***Index Terms*—generative AI, architectural education, design studio, creativity, critical thinking, academic integrity, design pedagogy**

I. INTRODUCTION

Architectural design education has always been built around a distinctive pedagogy: the studio, in which students iterate through sketches, models, and drawings under the close, often personal, supervision of a critic, and in which the artifact produced is simultaneously a technical exercise and an expression of individual authorship. Into this pedagogy has arrived a new class of tool. Text-to-image diffusion models, large language models (LLMs), and increasingly discipline-specific GenAI plug-ins can now generate massing studies, facade options, renderings, and even structural rationales in seconds, from little more than a short prompt.

This shift is not simply a faster version of familiar computer-aided design. Earlier digital tools — CAD, BIM, parametric scripting — extended the architect's hand; GenAI tools instead propose, on their own initiative, images and forms the student did not draw. That distinction matters pedagogically, because architectural education has historically treated the act of drawing as inseparable from the act of thinking through a problem. When a machine can generate a plausible design proposal before a student has fully formulated the question it answers, the traditional sequence of problem-framing, ideation, and representation is disrupted.

This paper reviews recent literature — much of it published in 2025 and 2026, reflecting how new this body of evidence is — across four themes central to evaluating GenAI's role in architectural education: student creativity, learning outcomes, critical thinking, and academic integrity. It then discusses what these findings imply for curriculum and studio pedagogy.

II. BACKGROUND: GENAI TOOLS IN THE ARCHITECTURE STUDIO

GenAI use in architecture schools spans a spectrum from generic consumer tools to bespoke, discipline-specific systems. Generic image generators are widely accessible and require no special training, but they lack architectural grounding — they do not reliably respect scale, structural logic, or building codes, and their aesthetic vocabulary is shaped by whatever imagery dominates their training data rather than by site or program. In response, researchers and educators have begun developing locally run, architecture-specific GenAI workflows intended to combine the accessibility of generic tools with greater contextual and technical relevance, while also aiming to support inclusive participation by students with weaker freehand drafting or rendering skills.

Systematic reviews of this literature describe GenAI as increasingly embedded across the design process — from pre-design analysis and concept generation through design development and production — while also cautioning that empirical evidence remains uneven across these stages, and that long-term effects on students' critical thinking and problem-solving are still poorly understood. This unevenness in the evidence base is itself a finding worth foregrounding before turning to what is known.

III. IMPACT ON STUDENT CREATIVITY

The most consistent finding across recent studies is that students experience GenAI as supportive of creativity, but unevenly across the design process. An exploratory study of architecture and interior architecture students found strong enthusiasm for GenAI image tools, which participants viewed as helpful during early-stage ideation and as boosting both creative confidence and perceived future professional competitiveness. Notably, this enthusiasm was conditional: students regarded AI use as ethically acceptable primarily when it was transparently acknowledged, and their comfort with it declined markedly in later design stages, where originality and independent critical judgment were seen as more central to the work.

A quasi-experimental study conducted at King Fahd University of Petroleum and Minerals compared a group of architecture students using GenAI tools in a supervised project against a control group working without them. The GenAI-assisted group generally outperformed the control group, with the clearest gains — an improvement of roughly fourteen percent — appearing in design form and aesthetic quality, alongside gains in time efficiency. The same study, drawing on interviews with practicing experts, proposed a framework recommending different levels of GenAI integration at different design stages rather than uniform use throughout.

Other work has tried to measure creativity more directly rather than relying only on self-report. One study built around a student-led AI drawing competition used a Creativity Support Index alongside semantic analysis of students' text prompts and aesthetic evaluation of the resulting images, finding meaningful variation in how creatively students formulated prompts and identifying "prompt literacy" — the skill of crafting generative instructions — as a distinct competency that design pedagogy will need to teach explicitly. A related study in a first-year architecture course structured GenAI use around Boden's theories of creativity through a three-phase sequence of combinational form generation, exploratory validation, and transformative prompt engineering, and found that self-evaluation played an important role in translating GenAI output into genuine design understanding rather than passive acceptance of machine-generated forms.

Taken together, this literature suggests GenAI's creative benefit is real but conditional: it appears to expand the range of forms students explore and to lower the barrier to visual experimentation, particularly for students with weaker manual drafting skills, but its value depends on how deliberately it is scaffolded, and its perceived legitimacy among students themselves narrows sharply once a project moves from ideation toward a defended, original design proposal.

.

IV. IMPACT ON LEARNING OUTCOMES

Evidence on formal learning outcomes is more limited than evidence on creativity, but the available studies point in a broadly positive direction for well-structured, supervised use. The KFUPM study's improvement in design form and aesthetics is one of the few outcome measures reported in architecture specifically. A broader meta-analysis spanning twenty-nine experimental and quasi-experimental studies across disciplines found that GenAI exerts a moderate positive effect on higher-order thinking overall, with the largest gains in problem-solving, somewhat smaller gains in critical thinking, and comparatively limited effects on creativity as a distinct outcome — a pattern that complicates any simple claim that GenAI straightforwardly improves creative learning.

That same meta-analysis identified an important moderating variable: duration of exposure. Interventions lasting up to roughly sixteen weeks tended to show the strongest gains, while longer periods of Gen-AI exposure were associated with diminishing returns, as sustained use fostered over-reliance, passive acceptance of AI output without critical examination, and a shift toward valuing efficiency over deep, self-directed inquiry. This finding is significant for architecture programs, where studio courses often run across a full semester or year and where the risk of habituation to AI-generated content is correspondingly higher than in shorter interventions.

This body of research also cautions against treating "learning outcomes" as a single variable. Gains in design form or rendering quality do not necessarily indicate gains in the underlying design reasoning that studio education is meant to cultivate — a distinction architectural educators will need to hold onto as they design assessment that credits process and judgment, not only the polish of a final image.

V. IMPACT ON CRITICAL THINKING

The critical thinking literature is where the strongest cautionary evidence appears. A systematic review synthesizing multiple studies on GenAI dialogue systems concluded that over-reliance on such tools is associated with a decline in critical and analytical thinking, particularly when students grow dependent on AI-generated content, with downstream effects including diminished creativity and reduced independent problem-solving capacity. Related work in vocational education has described this pattern as "metacognitive laziness" — a documented tendency for AI dependency to erode learners' self-regulation and critical engagement with material over time.

Mechanistically, researchers link this to how large language models present output: fluent, immediate, and authoritative-seeming, but without the argumentative scaffolding that would let a student assess why a given answer or form is correct. Scholars applying Cognitive Load Theory to this problem have argued that repeated reliance on generative tools as a substitute for effortful thinking gradually diminishes students' initiative and what they term epistemic agency — a particular risk for younger or less experienced learners who may not yet have the metacognitive strategies needed to critically evaluate AI content, and who may treat GenAI output as an authoritative conclusion rather than a starting proposal to be interrogated.

Not all findings are uniformly negative, however, and the evidence suggests that outcomes depend heavily on instructional design rather than on the technology itself. A study of critical thinking interventions embedded in problem-based learning found that explicit instruction — combining authentic tasks, structured dialogue, and AI-supported peer mentoring — did not significantly raise students' self-reported critical thinking in a short intervention window, but did measurably reduce thoughtless, unreflective use of GenAI and reduced students' tendency simply to adopt AI-generated content wholesale. Students in that intervention condition also produced more original, higher idea-density work than a comparison group, suggesting that structured critical engagement with AI output — rather than either unrestricted use or prohibition — is what protects and even strengthens creative and analytical performance. Complementary work has shown that critical thinking mediates how students rely on GenAI in problem-solving: students who actively assess AI outputs rather than accepting them passively are more likely to use these tools reflectively, cautiously, and collaboratively, while those with weaker critical evaluation habits are more prone to thoughtless reliance.

For architecture specifically, this points to a design-studio-appropriate strategy: treating GenAI output the way a good critic treats a first sketch — as a provisional proposal to be interrogated, defended, or discarded — rather than as a finished answer. Systematic reviews of GenAI in architectural education note this as an outstanding gap, calling explicitly for further research into GenAI's long-term effects on students' critical thinking and problem-solving skills across the full design process, not just at the ideation stage where most current studies concentrate.

VI. ACADEMIC INTEGRITY CHALLENGES

GenAI unsettles academic integrity frameworks that were built around a much simpler problem: uncredited copying of existing, identifiable work. Because LLMs and image generators produce output that is technically novel — assembled through probabilistic pattern-matching rather than direct duplication — that output resists the categories plagiarism policy was designed to catch. Scholars have coined the term "AI-giarism" to describe the resulting gray zone: students presenting AI-generated material as their own original work in ways that evade both students' own instincts about what counts as cheating and institutions' existing detection tools.

Survey research on student perceptions reinforces how unsettled this gray zone is. A 2025 survey of several hundred U.S. university students found that students' own ethical beliefs — not institutional policy — were the strongest predictor of whether they viewed AI-assisted work as misconduct and of whether they actually used AI tools this way; awareness of formal policy had no significant effect on either judgment or behavior. Many students in that study did not treat GenAI use as an extension of conventional cheating at all, but as belonging to a distinct, less clearly regulated category of conduct — a finding with direct implications for architecture schools, where studio culture already normalizes referencing, collaging, and reworking precedent images in ways that can make GenAI-assisted authorship feel like a natural extension of existing practice rather than a rule violation.

Detection is a further complication. Plagiarism-detection tools built for text struggle with AI-generated prose, and no equivalent detection infrastructure exists for AI-generated architectural imagery at all; a rendering or massing study produced substantially by a generative model leaves no reliable forensic trace the way copied text can. Faculty-focused research examining institutional policy in practice has found that instructors perceive existing academic integrity policies as poorly suited to "gray area" scenarios involving generative AI, and rate them significantly less favorably and less consistently than they rate traditional plagiarism policy — evidence that institutional frameworks generally, and architectural program policy specifically, are lagging behind the technology students are already using.

For architecture, the stakes are distinctive because studio evaluation has always rested on an implicit claim of authorship: a portfolio is presented, and graded, as evidence of what an individual student can design. Broader systematic reviews of AI and academic integrity call for institutions to develop clearer policy on the intellectual property status of AI-generated work and to redefine authorship for AI-assisted output more generally — a call architecture schools have particular reason to answer quickly, given how central authored visual work is to how the discipline credentials its students.

VII. DISCUSSION: TOWARD RESPONSIBLE INTEGRATION

Read together, this literature does not support either uncritical adoption or prohibition of GenAI in architectural education; it supports differentiation by design stage, paired with explicit instruction in critical evaluation and transparent disclosure norms.

7.1 Stage-Differentiated Use

The evidence consistently distinguishes early ideation, where GenAI appears to expand exploration and is broadly accepted by students, from later design development and production stages, where both students and the broader evidence base treat independent judgment as essential and where unchecked GenAI use carries the clearest risk to critical thinking. Curricula and studio briefs that explicitly name which stages invite open GenAI experimentation and which stages require the student to work substantially from their own developed reasoning give students a clearer, less ambiguous standard than a single blanket policy can.

7.2 Teaching Prompt Literacy and Critical Evaluation as Skills

Multiple studies identify "prompt literacy" — the ability to write generative instructions that meaningfully shape useful, architecturally coherent output — as a skill that itself requires teaching, not an incidental byproduct of access to a tool. Equally, the intervention research on critical thinking suggests that structured practices such as requiring students to justify, critique, or revise AI-generated proposals before incorporating them measurably reduce thoughtless adoption of AI output, even when they do not immediately raise self-reported critical thinking scores. Building such justification steps into studio review — asking not just what a student produced but why they kept, rejected, or reworked an AI-generated option — operationalizes this finding directly.

7.3 Transparency Over Prohibition

Because detection tools cannot reliably identify AI-generated architectural imagery, and because student research suggests ethical judgment is a stronger predictor of behaviour than policy awareness, disclosure-based norms are likely to be more enforceable and more educationally honest than detection-based enforcement. Requiring students to document which stages of a project involved GenAI assistance — treating this the way citation already treats textual sources — aligns academic integrity practice with how students in the reviewed studies already say they think about legitimate use.

7.4 Limitations of the Current Evidence Base

Much of the literature reviewed here is drawn from single-institution studies, short intervention windows, or fields adjacent to architecture rather than architecture itself, and several of the studies cited are recent preprints or newly published articles that have not yet accumulated a substantial body of replication. Systematic reviews specific to architectural education explicitly flag the absence of longitudinal evidence on critical thinking and problem-solving as a major gap. Conclusions in this paper should therefore be read as an evidence-informed synthesis of an early and rapidly evolving research area, not a settled verdict.

VIII. CONCLUSION

Generative AI has already changed how architecture students generate and explore form, and the early evidence suggests genuine creative and efficiency benefits, particularly in the ideation stages of design. But the same body of evidence warns that these benefits are conditional on deliberate pedagogical structure: unscaffolded, sustained use is associated with declines in critical thinking and independent design judgment, and architecture's studio culture of iteration and referencing leaves it exposed to authorship and integrity questions that current institutional policy is generally not equipped to answer. The path forward that emerges from this literature is not a choice between embracing or banning GenAI, but a more demanding task: differentiating its role by design stage, teaching its use as an explicit skill, and replacing unenforceable detection with honest disclosure — so that the tool augments the judgment architectural education exists to build, rather than substituting for it.

REFERENCES

[1] Huh, M., & Tracy, S. (2025). Students' perceptions of generative AI image tools in design education: Insights from architectural education. Education Sciences, 15(9), 1160. https://doi.org/10.3390/educsci15091160

[2] Rawan Alamasi, Omar S. Asfour, Riyad Ashmeel. (2025). The impact of generative AI on architectural design education: Insights from hands-on experience with architecture students. Ain Shams Engineering Journal (ScienceDirect). https://www.sciencedirect.com/science/article/pii/S2090447925006203

[3] Carlos Medel-Vera, Sandy Britton, William Francis Gates. (2025). An exploration of the role of generative AI in fostering creativity in architectural learning environments. Journal of Building Engineering / ScienceDirect. https://www.sciencedirect.com/science/article/pii/S2666920X25001419

[4] Sanghee Lee, Soomin Kang. (2025). The influence of generative AI with prompt engineering on creative design in architectural education. Journal of Asian Architecture and Building Engineering. https://www.tandfonline.com/doi/full/10.1080/13467581.2025.2552446

[5] Rawan Alamasi, Omar S. Asfour, Riyad Ashmeel. (2026). Applications of generative AI in architectural design education: A systematic review and future insights. Architecture, 6(1), 6. https://www.mdpi.com/2673-6470/6/1/6

[6] Aly, N., et al. (2026). Gen-AI-tecture: Using generative AI to support architectural students in design tasks. arXiv preprint. https://arxiv.org/pdf/2605.21361

[7] Brady Lund, Nishith Reddy Mannurru, Zoe abby Teel, Tae Hee Lee, Nathanlie Jugan Ortega, Sara Simmons, Evelyn Ward. (2025). Student perceptions of AI-assisted writing and academic integrity: Ethical concerns, academic misconduct, and use of generative AI in higher education.

(MDPI journal), 1(1), 2. https://www.mdpi.com/3042-8130/1/1/2

[8] Vəhid Gəruslu, Zafar Jafarov, Aytan Mövsümova, Leyla Məmmədova, Hüseyn Mirzayev. (2026). Implementing GenAI-supported learning in software engineering and computer science education using Bloom's Taxonomy. arXiv preprint. https://arxiv.org/pdf/2606.27398

[9] Rami Alsharefeen, Naji Al Sayari (2025). Examining academic integrity policy and practice in the era of AI: A case study of faculty perspectives. Frontiers in Education. https://www.frontiersin.org/journals/education/articles/10.3389/feduc.2025.1621743/full

[10] Himendra Balalle, Sachini Pannilage. (2025). Reassessing academic integrity in the age of AI: A systematic literature review on AI and academic integrity. (ScienceDirect). https://www.sciencedirect.com/science/article/pii/S2590291125000269

[11] Christine E. King, Beth A. Lopour. (2025). Fostering critical thinking during use of generative AI: A novel learning module for ideation in biomedical engineering design. Biomedical Engineering Education. https://link.springer.com/article/10.1007/s43683-025-00192-8

[12] Amir Yunus, Peng Rend Gay, Oon Teng Lee. (2026). From co-design to metacognitive laziness: Evaluating generative AI in vocational education. arXiv preprint. https://arxiv.org/pdf/2512.12306

[13] Mireia Vendrell, Samantha-Kaye Johnston. (2026). Scaffolding critical thinking with generative AI: Design principles for integrating large language models in higher education. (ScienceDirect). https://www.sciencedirect.com/science/article/pii/S2666920X26000342

[14] Chenyu Hou, Gaoxia Zhu, Yanzhi Liu, Vidya Su darshan, Josephine Leng Leng Chong, Fannie Yifan Zhang, Michael Yong Heng Tan, Yew Soon Ong. (2026). The effects of critical thinking intervention on reliance behaviors, problem-solving quality, and creativity during human-generative AI collaborative learning. (ScienceDirect). https://www.sciencedirect.com/science/article/pii/S036013152600014X

[15] Yan Zhao, Yuhe Yue, Zhonghua Sun, Qiang Jiang, Gangsheng Li. Does generative artificial intelligence improve students' higher-order thinking? A meta-analysis based on 29 experiments and quasi-experiments. PMC. https://pmc.ncbi.nlm.nih.gov/articles/PMC12734368/

[16] Hou, X., et al. (2025). The role of critical thinking on undergraduates' reliance behaviours on generative AI in problem-solving. British Journal of Educational Technology. https://bera-journals.onlinelibrary.wiley.com/doi/10.1111/bjet.13613

[17] Jinrui Tian, Ronghua Zhang. (2025). Learners' AI dependence and critical thinking: The psychological mechanism of fatigue and the social buffering role of AI literacy. (ScienceDirect). https://www.sciencedirect.com/science/article/pii/S0001691825010388