World Journal of Educational Studies
Print ISSN: 2959-9989
Online ISSN: 2959-9997
DOI: https://doi.org/10.61784/wjes3213

AI-DRIVEN, EVIDENCE-BASED GENERATIVE DESIGN: A TEACHING FRAMEWORK FOR LANDSCAPE ARCHITECTONICS INTEGRATING RESEARCH FINDINGS, DESIGN PRECEDENTS, AND INTELLIGENT RECOGNITION AND JUDGMENT

LiHao Yao, Rui Zhang, ZiXin Liang*
College of Coastal Agricultural Sciences, Guangdong Ocean University, Zhanjiang 524088, Guangdong, China.
**Corresponding Author: ZiXin Liang*

Abstract: Artificial intelligence (AI) has entered design courses, but its use often stops at tool operation. Taking Landscape Architectonics, an undergraduate course in landscape architecture programs, as the vehicle, this paper builds a teaching framework for generative design that integrates four elements: research findings, design precedents, AI-based recognition and judgment, and AI generation. The core claim is simple. Teaching students to "use AI" means teaching them to run an evidence-constrained generative design chain. "AI-driven" therefore carries two meanings: at the tool level, AI executes generation; at the teaching level, evidence steers generation, and the work of teaching lies in leading students from the former toward the latter. Among the four elements, research findings and design precedents are static knowledge that travels across sites. Research findings supply criteria for what counts as good. Design precedents supply examples of what to do, distilled from 13 street design guidelines into 114 retrievable situation packages. AI-based recognition and AI generation produce dynamic outputs anchored to a specific site. Students mobilize these elements through three stages—reading the site, grounding the context, and giving form—and a chain of evidence disciplines each stage. The framework follows three design principles: complementary elements, partial embedding, and explicit judgment. Three classroom examples show how the framework operates. This paper focuses on framework construction and does not report learning-outcome data.
Keywords: AI-driven generative design; Landscape architectonics; Evidence-based design; Curriculum reform; Evidence constraint

1 INTRODUCTION

The building industry is moving from expansion toward quality, and the knowledge structure of design work is changing with it. A scheme now has to defend itself with site data, environmental performance, and established research, not with formal intuition alone. Generative AI, meanwhile, has entered design workflows in substance. Design education has responded with a narrative centered on generation: teach students to write prompts, and let diffusion models produce scheme imagery quickly. This path eases the cold-start problem of a blank page. It also exposes a new problem. Once generation becomes cheap, the center of design must shift from production to judgment.

Landscape Architectonics is a core course in landscape architecture programs. It teaches the design of pavilions, covered walkways, and waterside pavilions, and it combines training in site cognition, climatic adaptation, and architectural expression. Three problems stand out in its current teaching. The first is subjective site cognition. Site surveying depends heavily on personal experience and on how well a student writes. Questions that could be quantified—how much of the ground is shaded by tree canopies, how much is hardened—are usually answered with phrases such as "well greened". Such statements cannot be checked, and they cannot support a design argument. The second is fragmented knowledge of environmental adaptation. Pavilions, walkways, and waterside pavilions are, at bottom, responses to climate. Yet sound, light, and heat are taught as isolated topics, disconnected from design decisions. Students can recite the points of passive design, but they cannot explain why a pavilion at a given location should be shaded in a given way. The third is the prioritization of expression over logic. Course assessment has long favored graphic quality, so students come to view design as producing attractive drawings. Generative tools, which now deliver polished images with ease, amplify this inversion.

Research on AI in architectural design courses over the past three years falls roughly into three lines: full-process intervention, coordination with design software, and multi-agent generation of teaching content. All three share a tendency. They make "letting AI generate" the main line of teaching and treat "judging whether AI is right" as an accessory. The cost of this choice has been recorded. One study measured generative efficiency and scheme logic at the same time: after AIGC was introduced, total design time fell by 51.0%, and scheme innovativeness and technical expression improved significantly, while functional-logic scores showed no significant gain [1]. Another study found that, without structured guidance, students used AI in a fragmented way [2]. A deeper problem is where correction is aimed. Existing studies correct AI output mostly at the image layer—stair junctions, slopes, structural support. The more dangerous errors sit at the diagnostic layer: AI claims to see what is not visible and states uncertain judgments as

© By the Author(s) 2026, under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0).

facts. Such errors do not appear in the drawing. Unless students are trained to check, they will take them as facts and design on that basis. The question of this paper is therefore not "how to make AI generate better", but "how to teach students to steer AI's generation".

The paper does two things. First, it builds a teaching framework made up of four elements (research findings, design precedents, AI-based recognition and judgment, AI generation), a three-stage process (reading the site, grounding the context, giving form), one chain of evidence, and four competence goals (Section 3). Second, it uses three classroom examples to show how the framework operates (Section 4). The paper focuses on framework construction. It does not report learning-outcome data; limitations and follow-up plans appear in Section 5.

2 RELATED WORK

Generative AI and design teaching. The field has formed a narrative centered on generation, and quasi-experimental work supports a positive role in early creative stages [3]. Correction of AI output, however, still targets image-layer errors, and most studies work with general-purpose generative models. Two recent empirical studies sharpen the problem. In the study that measured efficiency and logic together, introducing AIGC cut total design time by 51.0% and significantly improved innovativeness and technical expression, while functional-logic scores did not improve significantly [1]. AI applied to generation improves speed and presentation, not the quality of design reasoning. The second study found that students' AI use became fragmented when structured guidance was absent [2]. In landscape architecture, researchers have built an AI-driven landscape design teaching platform that supports design analysis with neural networks [4]. Others have validated an evidence-based design framework in an undergraduate urban design studio and found gains in how schemes were framed, developed, and evaluated [5]. Studies that make "judging AI output" itself a teaching objective remain few.

Evidence-based design and case-based teaching: two kinds of grounds for judgment. The two static knowledge elements of our framework correspond to these two literatures. Evidence-based design asks that design decisions rest on checkable evidence. In landscape architecture, it is formally defined as the deliberate and explicit use of scholarly evidence in land decisions [6]. A survey of architecture schools across Europe found that most curricula touched on evidence in some sense, but only a small minority named it as a curricular marker; the authors therefore recommend embedding evidence-based design in existing studios rather than offering separate courses [7]. Case-based teaching draws on the nature of design knowledge itself. Akin notes that knowledge in the studio exists mainly as precedents and heuristics, not as secure first principles [8]. Case-based reasoning in cognitive science gives a more general account: cases are the basic units of memory, reasoning, and transfer [9]. The two complement each other—evidence answers "what counts as good", and cases answer "what to do"—and early attempts to place them side by side exist [10]. Yet both share one risk: being treated as absolute standards. As Hamilton points out, evidence "supporting" a design is far from proof; evidence is rarely direct, and it often conflicts and shifts with context [11]. Our platform writes one rule into its constraints for exactly this reason: any threshold from any single study is a reference condition, not a universal standard. We found no work that brings the two into one teaching process as parallel knowledge sources, and none on the classroom use of AI-organized, AI-retrieved case libraries. That is where this paper sits.

3 FRAMEWORK CONSTRUCTION

3.1 Basis and Design Principles

The framework rests on three grounds: the teaching problems (subjective site cognition, fragmented environmental knowledge, and the prioritization of expression over logic; Section 1); the capability limits of AI (generation does not improve functional logic, and unstructured introduction leads to fragmented use; Section 2); and the constrainability of the carrier (if AI output is an opaque black box, evidence has nothing to attach to, and the framework has nowhere to land). Three design principles follow. Complementary elements: no single AI capability can support design judgment, because recognition yields site data and generation yields scheme images, and neither can say what counts as good; two kinds of professional knowledge, research findings and design precedents, must therefore enter. Partial embedding: AI is connected only where evidence can constrain it, because the applicability of AI is stage-dependent. Explicit judgment: every AI output must undergo explicit evaluation by a student before entering the design, and nothing passes by default. This turns "steering generation" from an attitude into an assessable act. Together, the three principles fix the shape of the framework. It is not a stack of tools. It is a structure organized around one question: how students, supported by evidence, can steer AI generation (Figure 1).

3.2 The Four Elements and Their Relations

Research findings are peer-reviewed conclusions that can serve as grounds for judgment. This element does not point to any single paper. Any study that yields testable criteria—thresholds, ranges, dominant factors, explanations of mechanisms—may enter. The course currently uses a body of work on the dominant factors linking urban form to heat resilience, with the study by Yao et al. as an example; comparable studies can replace or extend it [12]. The element supplies judgment rules. It answers "what counts as good", holds across sites and years, and is static knowledge. Design precedents are workable examples drawn from the platform's built-in knowledge base of street design guidelines and its

reference image library (construction and use in Section 3.3). This element supplies precedent methods. It answers "what to do", likewise crosses sites, and is likewise static.

The AI-based recognition and judgment element provides evidence that AI extracts from site imagery and anchors to the site itself: street-scene indicators from panoramic semantic segmentation, fisheye view factors, and, on top of these, the model's morphological typing and problem diagnosis. It supplies site evidence. It answers "what is needed here", varies with the site, and is a dynamic output. AI generation means design expression produced by AI under constraints, that is, controlled generation of concept renderings. It supplies scheme output. It answers "how to change this place", varies with the scheme, and is a dynamic output. The four are not four parallel technologies. They form a directed chain: static knowledge (research findings and design precedents) constrains dynamic output (recognition and generation). Without the first pair, the second can only judge by surface features of images and generate by formal preference. Without the second, the first stays inside journals and atlases and never reaches a real site. Their integration is the condition under which "AI-driven" holds.

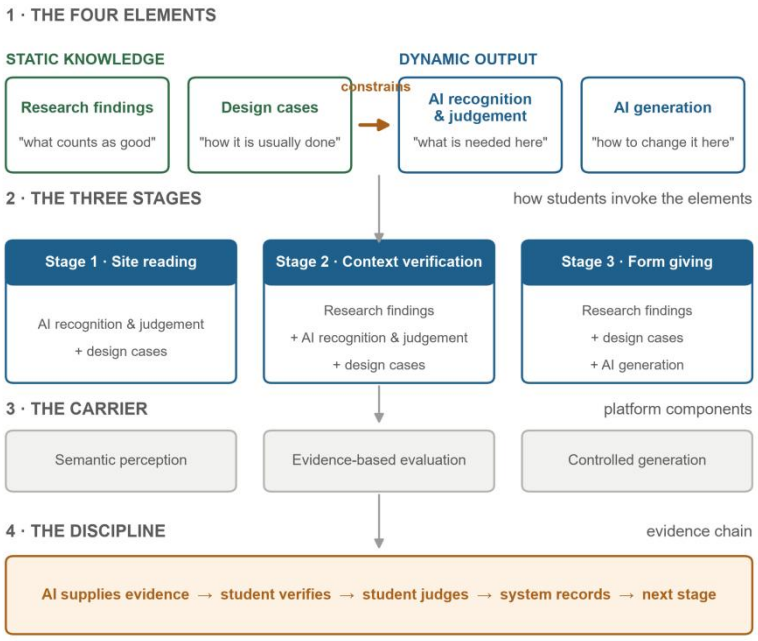


Figure 1 Framework Overview: Four Elements, Three Stages, the Platform Carrier, and the Chain of Evidence

3.3 Distilling Design Precedents, Classifying Problems, and Judging Them

From guidelines to a citable, structured case base. The precedent element is not a folder of exemplary pictures. It structures street-design knowledge from multiple cities into criteria that a model can cite. The platform holds 13 street design guidelines from China and abroad—London (2004), New York (2009), Shanghai (2016), Zhuzhou (2017), Shenzhen Luohu (2017), Beijing (2018), Guangzhou (2018, 2019), Yuxi (2018), Chengdu (2019), Foshan (2019), Qingdao (2020), Nanjing Qinhuai, and Qionglai—plus 19 interpretive papers. A six-step pipeline distills them into 114 "situation packages": block-level text extraction, chapter segmentation and quality checking, chapter distillation notes, element merging (large language model), package finalization (large language model), and validation and indexing. The image-processing pipeline runs in parallel. Illustrations are extracted from the 13 guidelines, matched with captions, deduplicated, and auto-tagged by element, yielding an indexed library of 3,869 images.

Three quality rules govern distillation. No harmonized numbers: city-specific values are listed side by side and labeled by city, and a "consensus value" is written only when at least three guidelines agree. Traceable sources: every point is annotated with a "guideline number–page" reference. Checkable images: each situation package must cite at least two city sources and three reference images that actually exist; packages that fail go into a "gap log" rather than being padded out. These rules separate the case base from an ordinary case collection: every piece of knowledge carries its source and its boundary of application.

Problem classification: 11 categories, 114 packages. The 114 situation packages are organized into 11 categories of street problems (Table 1). Each package contains nine fixed fields: applicable street type, recognition cues, guideline points (by city, with sources), recommended practices, consensus values, reference images, visual descriptors, negative prompt terms, and provenance.

Table 1 Eleven Categories of Street Problems and the Number of Situation Packages

Problem category	Packages	Problem category	Packages
Sidewalks and slow mobility	12	Building frontage, signage and outdoor display	14
Roadways and cross-sections	8	Paving and accessibility	8

Problem category	Packages	Problem category	Packages
Intersections and crossings	12	Lighting and smart facilities	7
Transit and rail connections	7	Parking and traffic organization	7
Greening and sponge measures	12	Overall character and color	13
Street furniture and municipal utilities	14	Total	114

The point of this structure is that a problem is no longer a vague "something is off here". It is a minimal knowledge unit with an ID, recognition cues, practices, and sources. It serves both as the candidate list during diagnosis and as the practice index during design.

Problem judgment: constrained multimodal diagnosis. For a given point, the platform submits the full panorama, the segmentation overlay, nearby points of interest, waterlogging records, and the road network to a multimodal model. The model must return structured results: the top two candidates among nine street types, and judgments covering the 11 problem categories, each with a location box on the image, visible evidence, a confidence level, and conditions still to be verified. Four "evidence gates" do the real work. Closed vocabulary: problem IDs must come from the 114 packages; the model may not invent problem types. Visible evidence: every problem must carry a location box and pictorial evidence, or it is rejected. Hazard downgrading: anything involving demolition, lane removal, or utility relocation is downgraded to a research suggestion and may not enter renderings. Full coverage: all 11 categories must be answered one by one, and "no problem seen" does not equal "no problem". Every automatic diagnosis is labeled a candidate. The user decides whether to accept, reject, or supplement each diagnosis.

Reference images: from precedent to practice. Every measure carries three things at generation time: the situation-package ID, city-level sources, and negative prompt terms. Reference images are retrieved from the 3,869-image library by package and scene; each package cites three to five site photographs across cities, with page numbers, and at most three are attached during generation. Their job is to guide materials and methods. They must not replace the buildings, roads, or spatial relations on site. Dimensions and practices from other cities are always marked as "reference", never as mandatory local standards—the same discipline as "thresholds are reference conditions" above. The precedent element thus shifts from "borrowing by looking at pictures" to "calling by ID and tracing by source". Students are not imitating an image. They are drawing evidence and methods from a knowledge system that has IDs, cues, practices, and boundaries.

3.4 Three Stages: How the Elements Are Used

The course's original teaching sequence had six phases: object basics, site cognition, design methods, form construction, environmental adaptation, and architectural expression. Under partial embedding, AI takes over only the three phases that evidence can constrain. They are regrouped into three stages—Site Reading, Context Grounding, and Form Giving—while the other three phases stay traditional. Table 2 lists the elements each stage uses and the competences it develops.

Table 2 Elements Used and Competences Developed in the Three Stages

Stage	Elements used	Student actions	Competences developed
Site Reading	AI recognition and judgment + design precedents	Verify segmentation, revise labels, derive site evidence	Can call, can constrain
Context Grounding	Research findings + AI recognition and judgment + design precedents	Verify problems item by item, interpret deviations	Can verify
Form Giving	Research findings + design precedents + AI generation	Write prompts, check generated images, record revisions	Can finalize

The two static knowledge elements run through all three stages; they are the constants. The two dynamic output elements divide labor across the stages; they are the variables. In Site Reading, site surveying shifts from producing qualitative descriptions to producing checkable evidence. Students must verify whether the segmentation matches the original images, revise labels with reasons, and judge which indicators are unreliable given image conditions. In Context Grounding, environmental knowledge turns into design grounds. Students check each item against dominant factors and synergistic reference ranges, make an explicit "present / absent / manually revised" judgment, and separate "verified deviation" from "dimensions that cannot be judged for lack of evidence". In Form Giving, the rendering stops being the endpoint and becomes an object to be tested. Students incorporate the evidence from the first two stages into prompts, check the generated image against the original—massing, doors and windows, signage, entrances, road edges—and record every generation error they find. Across the three stages, part of the weight of design assessment moves from graphic quality to judgment quality.

3.5 The Underlying Mechanism: Chain of Evidence and Division of Responsibilities

All three stages share one chain: AI proposes evidence, students verify it, students make the judgment, the system keeps the trace, and the next stage begins. The chain does not restrict generation; it makes generation trustworthy. An unverified AI result carries no design weight, and students learn nothing of "steering". Once evidence supports a

generated design, a pretty picture becomes a defensible scheme. The chain of evidence is the passage by which "AI-driven" moves from the tool layer to the teaching layer.

The chain is, in substance, a division of responsibilities. AI proposes testable judgments. Students confirm whether each holds. Teachers decide whether a judgment enters the design. None of the three can be replaced. AI bears no responsibility for engineering feasibility, property rights, cost, or structural safety. Students may not take AI output as a conclusion. Teachers may not waive review because "the system has spoken". This division matches AI-literacy frameworks for design students that list "ethics and critical judgment" as a standalone dimension [13]. Using AI critically is itself a competence that can be developed and should be assessed (Figure 2).

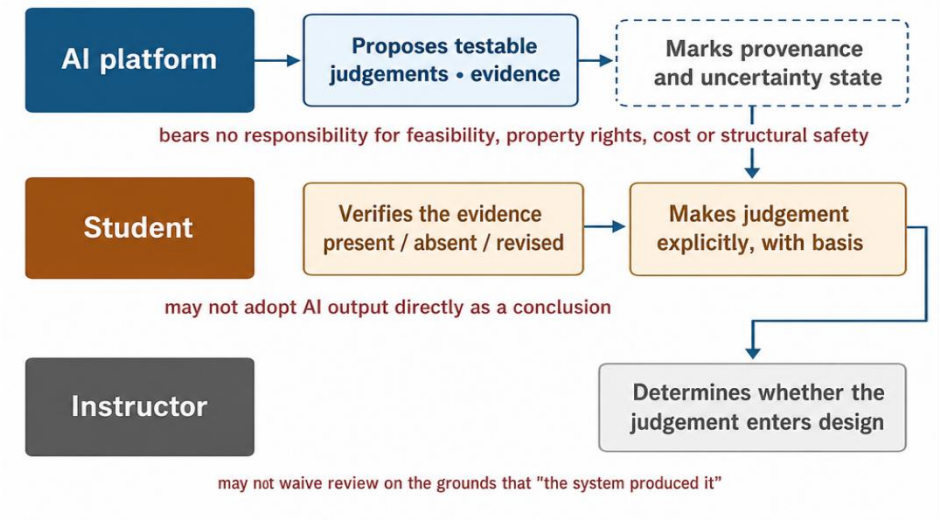


Figure 2 The Chain of Evidence and the Division of Responsibilities

3.6 Two Course-Specific Interfaces

First, view factors are the quantitative language of pavilions, walkways, and waterside pavilions. The design of buildings for rest is, at bottom, a balance between shade and openness, and the sky, tree-canopy, and building view factors describe exactly these three quantities. The platform is therefore more than a general analysis tool: it enters the design discussion of garden buildings themselves. A student can argue, for instance, that adding a section of covered walkway lowers the sky view factor by a given amount and thereby improves shade. Second, garden buildings are climatic devices. Pavilions, walkways, and waterside pavilions are themselves responses to climate. The heat-and-waterlogging diagnosis from Context Grounding supplies the evidence for questions such as why a pavilion at this spot should be shaded this way, or why this walkway should be ventilated that way. Climate-adaptation knowledge that once stayed at the level of experience becomes a defensible basis for design.

3.7 The Carrier: Platform and Constraint Mechanisms

The framework runs on a self-developed evidence-based AI platform. The platform is not teaching software in itself; part of this work is to show how a tool built for professional practice can be reshaped pedagogically for the classroom. It has three components: semantic perception (panoramic semantic segmentation and fisheye view factors); evidence-based evaluation (reading planimetric indicators on a grid, merging them with street-scene indicators into 13 items, and outputting morphological typing, dominant factors, and synergistic reference ranges); and controlled generation (producing concept renderings based on a fixed original image, prompt, and scheme version).

Three constraints turn the components into a teaching device. Explicit pending verification: every diagnosis defaults to "to be verified", must be confirmed item by item as "present / absent / manually revised", and cannot be submitted without a written rationale. Fixed constraints: real-world dimensions in meters may not be inferred from pixels; a single daytime image cannot determine nighttime illuminance, vehicle speed, utilities, or drainage capacity; and not seeing a facility does not mean it is absent. Separation of responsibilities: high-impact engineering conditions—demolition, lane removal, utility relocation—can only be marked for further study and may not enter design measures directly. Together, these constraints make AI's uncertainty explicit and require it to be addressed, rather than allowing it to remain a hidden defect. If AI output were an opaque black box, research findings and design precedents would have no purchase, and students would have nothing to steer. That is the technical premise on which the four-element integration stands.

4 CLASSROOM APPLICATION OF THE FRAMEWORK

The three examples below come from real platform runs, and their data entries and generation records are traceable. They are classroom demonstrations of student work, used to show how the framework operates and how the four elements are used. All three unfold on the same site—point 2158 in the 150 m analysis grid—and correspond to the

three stages in turn. Their organization is itself a demonstration of the framework: the evidence produced in one stage becomes the input to the next (Figure 3).

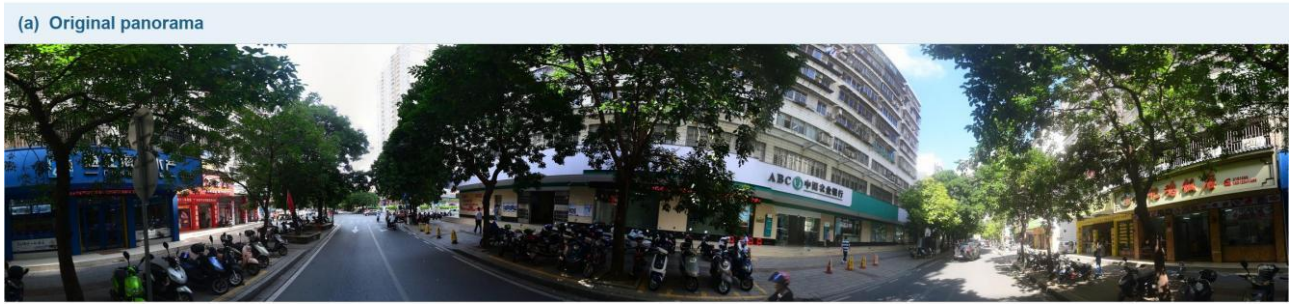


Figure 3 The Case Site: The Original Panorama

In Site Reading, students verify the semantic segmentation and the 13 indicators. The teaching point is a tension between measurement bases. At the street-scene scale, the tree-canopy share is 0.435 and the tree-canopy view factor is 0.385, so shading is good. This stage teaches students to ask about the measurement basis before interpreting the number. In Context Grounding, students verify each item against dominant factors and synergistic reference ranges. The teaching point is a tension of judgment. At this point, the dominant factor for high heat (HH) is the planimetric impervious-surface ratio, 0.992, beyond its reference range (< 0.86), while the dominant factor for heat resilience (HR) is the building view factor, 0.044, within its range (< 0.20). Students cannot write the site off as "a poor environment". They must separate "hardening should be controlled" from "building enclosure is acceptable". This stage teaches students to distinguish between specific findings rather than generalize. In Form Giving, students incorporate the evidence from the first two stages into prompts, generate renderings, and check them item by item. The teaching point comes from two real review records left by the platform. In the first, the initial comprehensive-renovation rendering mistakenly blocked an upstairs window, and this error was corrected through targeted revision. The second record states: "the canvas contains white margins; no stretched registration is performed; actual cooling is not verified". The first is a graphic deviation—something was drawn wrong, and revision fixed it. The second is insufficient evidence—the point is not where the drawing is wrong, but what the drawing cannot prove. Traditional teaching trains students to address the first and barely touches on the second. This stage trains students to tell graphic deviation from insufficient evidence. Together, the three examples show the order in which the four elements are used: the first uses AI recognition and design precedents, the second adds research findings, and the third adds AI generation, where all four converge. A student who finishes the three examples has completed one full pass from executing generation to steering it (Figure 4).



Figure 4 Detecting and Revising Generation Errors

5 DISCUSSION

The first limitation is that this paper reports no learning-outcome data. The study focuses mainly on framework construction; no large-scale pilot has been conducted, so we make no claims about teaching effects, and a controlled pilot under current conditions is the next step. Second, the platform's reach is limited. Its knowledge core comes from a study of one city's built-up area, so its morphological classes and synergistic thresholds are regional and need recalibration when transplanted. View factors are reference estimates that depend on projection assumptions and image quality; they are used for scheme comparison, not field measurement. These limits should themselves be taught, not hidden. Third, AI hallucination cannot be eliminated. Fixed constraints and explicit pending verification reduce visible harm, but no diagnosis is guaranteed to be correct.

The framework's transferability comes from its method, not its object. Evidence constraints, explicit pending verification, judgment traces, and the division of responsibilities do not depend on urban streets. To move to park green space, campus landscapes, or traditional settlements, one replaces the site data sources and evaluation indicators; the three stages and the chain of evidence carry over. But transfer presupposes a constrainable tool. With a general-purpose generative model, "evidence constraint" cannot be implemented. That is the fundamental reason this paper chooses a self-developed domain platform over off-the-shelf tools.

6 CONCLUSION

Now that generative AI has entered design workflows, the core question for design education is not whether to use AI but how to teach students to steer its generation. Taking the Landscape Architectonics course as the vehicle, this paper has built a teaching framework for generative design that integrates four elements: research findings, design precedents, AI-based recognition and judgment, and AI generation. The main line is this: to teach students to "use AI" is to teach them to run an evidence-constrained generative design chain. "AI-driven" has two layers. At the tool layer, AI executes generation; at the teaching layer, evidence steers generation, and the whole effort of teaching lies in leading students from the former to the latter. Future work will proceed in two directions: conducting a controlled teaching pilot under current conditions to obtain learning-outcome data and testing whether the framework transfers to other planning and design courses.

COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

FUNDING

This research was supported by the Undergraduate Education and Teaching Reform Project of Guangdong Ocean University (PX-1852026004) and the Scientific Research Startup Funds of Guangdong Ocean University.

REFERENCES

[1] Xiangyu Z, Guoqing L, Li S. Reconstruction of Digital Design Teaching and Paths of Ideological and Political Integration in Landscape Architecture Driven by AIGC. Journal of Landscape Research, 2026, 18(1/2): 54-59.
[2] Jin S, Tu H, Li J, et al. Enhancing architectural education through artificial intelligence: a case study of an AI-assisted architectural programming and design course. Buildings, 2024, 14(6): 1613.
[3] Chandrasekera T, Hosseini Z, Perera U. Can artificial intelligence support creativity in early design processes? International Journal of Architectural Computing, 2025, 23(1): 122-136.
[4] Yang X, Jiang D, Liu M. Construction and application of a landscape design teaching platform driven by artificial intelligence. International Journal of Web-Based Learning and Teaching Technologies (IJWLTT), 2024, 19(1): 1-16.
[5] Shi S. Effectiveness of evidence-based design approaches for built environment professional education: A case study on an urban design studio. Buildings, 2024, 14(3): 836.
[6] Brown R D, Corry R C. Evidence-based landscape architecture: The maturing of a profession. Landscape and Urban Planning, 2011, 100(4): 327-329.
[7] Nikolić O, Careva K, Vujović M. Evidence-based design in architectural education: Curriculum visibility, recognition and interest. Facta Universitatis, Series: Architecture and Civil Engineering, 2026.
[8] Akin Ö. Case-based instruction strategies in architecture. Design Studies, 2002, 23(4): 407-431.
[9] Kolodner J L, Cox M T, González-Calero P A. Case-based reasoning-inspired approaches to education. The Knowledge Engineering Review, 2005, 20(3): 299-303.
[10] Pilosof N P, Grobman Y J. Evidence-based design in architectural education: Designing the first Maggie's Centre in Israel. HERD: Health Environments Research & Design Journal, 2021, 14(4): 114-129.
[11] Hamilton D K. There is evidence to support design for health, but rarely proof. HERD: Health Environments Research & Design Journal, 2018, 11(2): 189-191.
[12] Yao L, Zhang R, Wu J, et al. Heterogeneous dominant factors of heat resilience across diverse urban patterns. Land Use Policy, 2025, 157: 107676.
[13] Lin S, Zhengtang T. AI Literacy Scale for Undergraduate Design Students: Development and Validation. SAGE Open, 2026, 16(2): 21582440261463131.