

Learning to See with Machines: Aesthetic imagination and AI in architectural education

Victor Sardenberg¹, Nieri Soares de Araujo¹, Rafael Perrone¹

¹Universidade Presbiteriana Mackenzie, PPGAU, São Paulo, Brazil

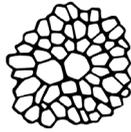
victor.sardenberg@mackenzista.com.br; nieri.araujo@mackenzie.br;
rafaelantonio.perrone@mackenzie.br

Abstract. This study investigates the integration of artificial intelligence in architectural education through two workshop editions using ComfyUI, a node-based interface for Stable Diffusion. Conducted in the second semester of 2024 and the first semester of 2025, the workshops engaged architecture students in generating images from textual prompts and critically reflecting on the outcomes. Combining hands-on experimentation with structured feedback, the study explores how AI-generated imagery influences students' design thinking, visual literacy, and understanding of authorship. Quantitative and qualitative data were collected through questionnaires and discussions, revealing high engagement, creativity stimulation, and a rapid learning curve—despite technical and conceptual challenges. By comparing both cohorts, the research identifies a pedagogical evolution from initial excitement toward more critical and contextualized uses of AI tools. The paper contributes to the discourse on design education by proposing a framework for incorporating generative AI into curricula, emphasizing reflective use, aesthetic exploration, and the development of new representational literacies.

Keywords: Artificial Intelligence, Architectural Education, Generative Design, Visual Literacy, Computational Aesthetics

1 Introduction

In recent years, artificial intelligence (AI) has emerged as a transformative force in architectural education and practice. Generative models, such as Stable Diffusion, have introduced new possibilities for visual ideation, enabling students to explore spatial concepts rapidly through image synthesis. Within this context, prompt-based image generation has gained traction as an



accessible design strategy that bridges textual imagination and visual output, fostering novel workflows for architectural experimentation.

Among the available tools, ComfyUI, an open-source, node-based interface for Stable Diffusion, has proven especially relevant for educational environments. Its modular structure and visual programming paradigm make it approachable for students with varying levels of technical expertise, while its image generation capabilities enable quick iteration and conceptual expansion. As architectural pedagogy increasingly embraces digital and speculative methods, tools like ComfyUI offer a way to stimulate creativity, provoke critical reflection, and reframe architectural representation beyond traditional photorealism.

This paper presents the findings of two experimental workshops conducted in the second semester of 2024 and the first semester of 2025, which integrated ComfyUI into undergraduate architectural education. Each edition engaged approximately 10 students in prompt-driven image generation exercises, followed by structured group discussions and critical feedback sessions. The aim was to investigate how AI-generated imagery affects students' design thinking, visual literacy, and attitudes toward authorship and creativity.

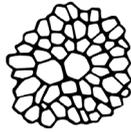
By comparing the outcomes of both cohorts, this study aims to understand how the pedagogical integration of AI has evolved over time and to what extent students have developed more critical, nuanced, or speculative uses of the tool. The central research questions guiding this investigation are:

- (1) How does prompt-based image generation with AI reshape the process of architectural learning?
- (2) What pedagogical transformations and challenges emerge when such tools are applied across different semesters and student groups?

Through this comparative analysis, the paper contributes to ongoing debates about the role of generative AI in architectural education, highlighting both its potential and its tensions as a catalyst for rethinking design methodologies, visual practices, and curricular structures.

2 Theoretical Framework

The integration of generative AI into architectural education intersects with multiple domains, including design pedagogy, computational aesthetics, and media theory. To understand the pedagogical implications of tools like ComfyUI, this study draws on three interconnected perspectives: visual literacy in design education, computational co-authorship, and aesthetic theory in the age of synthetic images.



2.1 Visual Literacy and Pedagogical Shifts

Architectural education has long relied on the development of representational skills to communicate spatial ideas. With the advent of AI-generated imagery, the ability to read, critique, and compose visual content becomes increasingly complex. The workshops in this study are grounded in pedagogical theories of experiential and constructivist learning (Schön, 1983; Kolb, 2014), where knowledge emerges from the process of doing, reflecting, and iterating. Stable diffusion becomes not just a design tool but a medium for engaging students in speculative and reflective thinking.

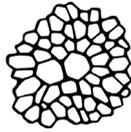
2.2 AI and the Question of Authorship

The use of AI in design challenges traditional notions of authorship. As students prompt a machine to generate images, agency is distributed between humans and algorithms. This agency aligns with recent theoretical work on computational co-authorship (Carpo, 2023), in which algorithms function not only as tools but as collaborators. The workshops provided a structured environment for students to navigate this shift by dynamically testing, correcting, and reinterpreting the machine's outputs.

2.3 From Photorealism to New Aesthetic Possibilities

Aesthetic discourse in architecture has historically been shaped by dominant visual styles, most recently exemplified by photorealistic renderings generated by software such as Lumion or Enscape. As Michael Young (Young, 2021) argues, these images often reinforce a narrow band of what architecture “should” look like. Generative AI, in contrast, opens a speculative space for reimagining architectural form and atmosphere. By engaging with ComfyUI and prompt-based generation, students explore images of architecture that deviate from normative representations and enter new aesthetic territories (Sardenberg & Perrone, 2025b; Sardenberg & Armagno Gentile, 2024; del Campo & Leach, 2022; May, 2017).

Together, these frameworks inform the design and analysis of the workshops. They provide the basis for evaluating how students learn not only to use AI tools, but to critically question their outputs—an essential skill in an era where architectural images are increasingly mediated by synthetic intelligence.



3 Methodology

This study is based on a comparative analysis of two editions of an experimental workshop on artificial intelligence in architectural education, conducted in the second semester of 2024 and the first semester of 2025. Both workshops were implemented as an optional seminar on generative AI at Universidade Presbiteriana Mackenzie and involved approximately 10 undergraduate architecture students per semester.

3.1 General Structure

Each edition was structured over multiple weekly sessions (6 to 8 meetings, each lasting 1 hour and 30 minutes), combining technical instruction, hands-on experimentation with Stable Diffusion via ComfyUI, and critical discussion. All students worked with the same software stack—ComfyUI with ControlNet and IPAdapter, and ChatGPT 4o for prompt generation—running on lab or personal machines. The students were initially encouraged to generate a few dozen lower-resolution images (approximately 1024 x 768), select a few to reuse in the iteration process, and then choose two or three to upscale four times to produce the final image.

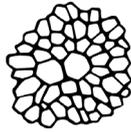
The workshops aimed not only to teach tool operation but also to critically reflect on how generative AI transforms representation, authorship, and design imagination in architectural contexts. Students were invited to iteratively produce, select, and present images in response to a series of guided design exercises. The aspect of quickly producing dozens of images was emphasized so that students should not follow their first intuition, but instead get lost and discover new design directions in the process by collaborating with the generative models.

3.2 Design Exercises

Three core exercises were conducted in both editions, progressively increasing in interpretive openness:

Exercise 1 – Architectural Canon in Cinematic Styles: Students were asked to represent canonical modern buildings (e.g., Villa Savoye, Casa da Cascata) using cinematic visual styles (e.g., Blade Runner, Metropolis, Wes Anderson). The goal was to examine how aesthetic codes from cinema could reframe familiar architecture, questioning representational norms and triggering discussions about genre, atmosphere, and narrative.

Exercise 2 – Imagining the Minhocão: The second exercise explored speculative futures for the Minhocão, a controversial elevated highway in São Paulo. Students generated interventions that were either utopian, critical, or surgical, using prompt engineering and iterative refinement. This broadness encouraged them to project social, ecological, and formal ideas into the city's fabric using AI imagery.



Exercise 3 – AI in the Design Studio: In the final stage, students were invited to apply Stable Diffusion to their architectural design proposals, which they had developed in parallel within their studio courses. This open-ended assignment aimed to assess the real-world applicability of AI tools for conceptual development, visualization, and exploration of design alternatives.

3.3 Data Collection

The research employed a mixed-methods approach, collecting both quantitative and qualitative data across the two editions:

Questionnaires: A structured survey with three parts:

- (1) participant background,
- (2) Likert-scale evaluation (ease of use, creativity, aesthetics, etc.),
- (3) open-ended reflections on the tool's impact.

Image Analysis Discussions: Visual outputs from each exercise were presented on the beamer and analyzed by the tutors and all participants for thematic diversity, aesthetic direction, student intent, challenges, and solutions.

Final Presentations: Students were asked to present a selection of generated images, explaining their process, aesthetic decisions, and design intentions.

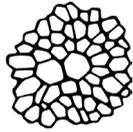
4 Results

The workshops generated a diverse range of image outputs, reflections, and student feedback, providing insights into how generative AI tools like ComfyUI are received, understood, and applied within the context of architectural education. This section presents findings structured around the three core exercises and contrasts responses across the two workshop editions (2024 and 2025).

4.1 Exercise 1 – Architectural Canon in Cinematic Styles

Students reimagined canonical modernist buildings in the visual styles of specific films (Figure 1). Outputs ranged from photorealistic reinterpretations (e.g., Casa da Cascata as seen through Blade Runner) to more surreal or stylized versions (Villa Savoye rendered in Wings of Desire aesthetics).

Common themes included exaggerated atmosphere, color grading, and compositional framing, revealing a newfound sensitivity to image genre and mood. Students reported increased awareness of how representational codes frame architectural meaning. Between editions, a noticeable shift occurred: while the 2024 cohort often focused on “getting the style right,” the 2025 cohort more consciously subverted styles, blending genres and reflecting on the cultural implications of aesthetic choices.



NY Guggenheim + 2001



Heydar Aliyev Center + Dune



Ronchamp + Brazil



Villa Savoye + Bridgerton



Robie house + Mad Max



Casa das Canoas + Blade Runner

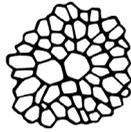


NY Guggenheim + Metropolis



Villa Savoye + Wings of Desire

Figure 1. Images of canonical buildings are reimagined with the aesthetics of films and series. Source: Authors, 2025.



4.2 Exercise 2 – Imagining the Minhocão

The speculative redesign of the Minhocão elicited a highly diverse range of interpretations (Figure 2). In 2024, students tended to create images as provocative visual statements, such as utopian parks, surreal megastructures, and dystopian ruins. In 2025, students, likely benefiting from refined prompt-writing skills and increased technical familiarity, articulated more grounded proposals, often combining surreal aesthetics with coherent urban narratives. In both cases, the use of AI provided a rapid means of testing alternative visions, bypassing the constraints of traditional representation techniques and allowing more speculative thought to emerge.

Across both editions, students emphasized how AI enabled them to explore futures that extended beyond conventional renderings or master plans. Recurring themes included verticalization, green reappropriation, infrastructural transformation, and symbolic interventions. Some proposals approached satire, others attempted realism through stylized compositions. The task stimulated engagement with pressing urban questions—mobility, public space, ecology—but filtered through a generative, image-based methodology that foregrounded visual experimentation as a mode of critical inquiry.

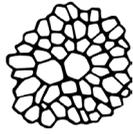
4.3 Exercise 3 – AI in the Design Studio

In the final exercise, students integrated ComfyUI into their ongoing studio projects. Many used AI to test formal languages, simulate materiality, or envision atmosphere prior to traditional modeling or drafting.

Students with less prior graphic experience reported that image generation helped them “see” their project more clearly and earlier in the process. Others cautioned about over-reliance on AI’s seductive visuals and stressed the need for critical interpretation.

A few students developed workflows that combined hand sketches, diagrammatic overlays, and AI-generated backgrounds, demonstrating hybrid approaches to design.

Notably, the 2025 cohort demonstrated more advanced image editing and compositing techniques (e.g., prompt travelling, seed manipulation, ControlNet-based corrections), indicating an increase in proficiency with the tool over time. These results indicate growing enthusiasm and ease with the tool, as well as a clear demand for curricular integration.



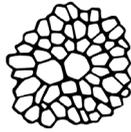
SIGraDi 2025 | Meta - Responsive Approaches



faud | UNC



Figure 2. Reimaginings of the Minhocão. Source: Authors, 2025.



4.4 Qualitative Themes

Analysis of open-ended survey responses revealed several recurring insights:

- Empowerment through iteration: Students appreciated being able to explore many ideas quickly.
- Surprise and inspiration: Unexpected outputs triggered new design directions.
- Tool limitations: Several noted challenges with scale accuracy, architectural logic, and prompt specificity.
- Critical distance: Some voiced concern over overreliance and the seductiveness of “good-looking” images.

5 Discussion

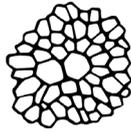
The integration of generative AI into architectural education—particularly through tools like ComfyUI—reveals not only new technical possibilities but also profound shifts in how students conceptualize authorship, design thinking, and the role of representation. Across the two workshop editions, a clear evolution emerged: from initial curiosity and playful experimentation to increasingly critical engagement and creative control.

5.1 AI as Co-Author: Redefining Design Agency

One of the most significant pedagogical shifts observed was the redistribution of authorship. Rather than viewing AI as a passive tool, students are increasingly engaging with it as a collaborative agent—a synthetic partner whose outputs can surprise, challenge, or expand their intentions. It is interesting to notice that at the beginning, they are very uncomfortable with the generative model not outputting exactly what they had in mind. In a second moment, they realize that it is necessary to enter a game between the designer and the generative model to explore and refine designs.

5.2 Visual Literacy and Design Imagination

Students reported that the rapid generation of diverse visual outputs enhanced their ability to “see” the project before modeling or drawing it. This anticipatory visioning, unbound by technical constraints, supported more speculative and open-ended design exploration. Importantly, the exercises sharpened not only their aesthetic sensitivity but also their critical visual literacy—the capacity to analyze, question, and edit synthetic imagery. In doing so, students began to recognize image generation not just as a visualization step, but as a design-thinking process in its own right. Moreover, they developed a vocabulary to



discuss aesthetics, which is uncommon in the modern curriculum of architectural education.

5.3 Pedagogical Potentials and Tensions

Three significant pedagogical contributions emerged from these workshops:

- Acceleration of design iteration, allowing for fast testing of alternatives;
- Aesthetic expansion, beyond the default photorealistic rendering paradigm;
- Critical reflection, prompting questions around authorship, ethics, and visual culture.

Nevertheless, the workshops also revealed key tensions: technical barriers to access, ambiguity in scale and spatial logic, and the risk of seduction by “beautiful” images that are disconnected from architectural coherence. These tensions reaffirm that AI in architectural education cannot be framed purely as a technical skill; instead, it must be embedded within discursive, reflective, and collaborative pedagogies.

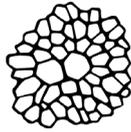
5.4 Comparative Insights: 2024 vs. 2025

Comparing the two cohorts reveals a trajectory of increasing confidence and depth: The 2024 edition emphasized accessibility, surprise, and the novelty of the AI tool itself. The 2025 edition, benefiting from accumulated experience and improved instruction, demonstrated greater conceptual maturity: prompts were more deliberate, aesthetics were more targeted, and critiques were more grounded. This pedagogical evolution suggests that single-semester exposure is insufficient. Real impact comes with continuity—AI tools must be embedded across curricular moments to move from surface-level effects to deep, transformative learning.

5.5 Future Work

The workshops served as a testing ground for a broader pedagogical and research agenda. Building on the promising results, future work will pursue four key directions:

Computational Aesthetic Evaluation: Applying a computational aesthetics framework (Sardenberg, 2024) to assist students in exploring the design space of generative models. While the workshops foregrounded creative exploration, they also generated a rich visual dataset—hundreds of student-generated images that reflect emerging patterns, biases, and stylistic tendencies. The next step is to apply models such as the PHR (Sardenberg & Perrone, 2025a), which quantifies aesthetic features like visual complexity, novelty, and formal coherence. By mapping this data, the goal is to reveal how student preferences evolve, how aesthetic conventions are challenged or reinforced by AI, and how



computational analysis can support—not replace—design exploration and critique.

Ongoing Seminar Integration: Establishing a permanent thematic seminar in the curriculum where AI tools and critical theory co-evolve, allowing students to engage with technology not only as users but as theorists and experimenters.

Public Exhibition: Organizing an exhibition of student works that juxtaposes speculative, cinematic, and project-based outputs from both workshop editions, inviting public debate about the aesthetics of the artificial and the role of imagination in architectural education.

5.6 Reframing the Role of AI in Architectural Education

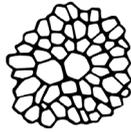
What is at stake is not whether AI will enter architectural education—it has already done so. The question is how its presence will be shaped: as a shortcut to generic images or as a tool for radical imagination. The workshops described here suggest that, with appropriate pedagogical framing, AI can do more than decorate a design—it can disrupt assumptions, catalyze reflection, and amplify the plurality of what architecture might become.

This pedagogical engagement with AI intersects directly with broader questions concerning the aesthetic and political dimensions of architectural education. Aesthetics, far from being a secondary concern, constitutes a fundamental mode through which architecture negotiates meaning, agency, and cultural relevance. In line with Jacques Rancière's concept of the distribution of the sensible (Rancière, 2013), aesthetics structures the field of visibility and intelligibility, determining what can be seen, imagined, or acted upon. To design is also to take a position on which worlds are possible, desirable, or even imaginable.

The outputs generated in these workshops—whether speculative visions for the Minhocão or cinematic reinterpretations of canonical buildings—serve as more than visual exercises. They represent aesthetic propositions that participate in the ongoing construction of architectural discourse. Through generative AI, students engage in aesthetic world-making, where computational tools are not merely assistive but constitutive of new modes of seeing, thinking, and designing.

Such an approach demands a reframing of how design is both taught and evaluated. Generative models like Stable Diffusion offer access not only to alternative aesthetic categories but to alternative epistemologies—ones that resist normative photorealism and encourage critical reflection on representation itself. The question thus becomes: What kinds of futures do these images make thinkable? What desires and ideologies are encoded in their aesthetics? What remains unrepresented or unrepresentable within the system?

Architectural education, when combined with generative AI and critical aesthetic inquiry, has the potential to cultivate designers who are not only image-makers but curators of cultural imagination. The objective is no longer to



train students merely to draw buildings, but to construct futures—to communicate with machines, ask more effective questions, and translate abstract intentions into situated images. In a moment when algorithmic systems increasingly mediate architecture, this capacity to imagine otherwise may become its most vital task.

Acknowledgements. This research is funded by CNPq (National Council for Scientific and Technological Development) and CAPES (Coordination for the Improvement of Higher Education Personnel) through the Institutional Program for Postdoctoral Research. Their support is gratefully acknowledged.

References

- Carmo, M. (2023). *Beyond Digital: Design and Automation at the End of Modernity*. The MIT Press.
- del Campo, M., & Leach, N. (2022). Can Machines Hallucinate Architecture? AI as Design Method. *Architectural Design*, 92(3), 6–13. <https://doi.org/10.1002/ad.2807>
- Kolb, D. A. (2014). *Experiential Learning: Experience as the Source of Learning and Development*. FT Press.
- May, J. (2017). Everything Is Already an Image. *Log*, 40, 9–26.
- Rancière, J. (2013). *The Politics of Aesthetics* (G. Rockhill, Trans.; Reprint edition). Bloomsbury Academic.
- Sardenberg, V. (2024). *Computational aesthetics in architecture: A framework for quantifying preferences using computer vision and artificial neural networks* [Doctoral Thesis, Hannover: Institutionelles Repositorium der Leibniz Universität Hannover]. <https://doi.org/10.15488/17879>
- Sardenberg, V., & Armagno Gentile, Á. (2024). *Silicon-Based Imagination of the “Invisible Cities”—How far can artificial intelligence hallucinate?* SIGraDi 2024, Barcelona.
- Sardenberg, V., & Perrone, R. (2025a). Expanding Design Creativity with the PHR2 Model: Predicting Hedonic Responses in Architecture. *Computational Design and Computer-Aided Creativity Workshop in the International Conference for Computational Creativity*. <https://computationalcreativity.net/workshops/computational-design-iccc25/papers/sardenberg2025a.pdf>
- Sardenberg, V., & Perrone, R. (2025b). *Navigating Architectural Aesthetics Beyond the Beautiful: Data-Driven Approaches to Design Evaluation* (SSRN Scholarly Paper No. 5134280). Social Science Research Network. <https://doi.org/10.2139/ssrn.5134280>
- Schön, D. A. (1983). *The Reflective Practitioner: How Professionals Think In Action*. Basic Books.
- Young, M. (2021). *Reality Modeled After Images: Architecture and Aesthetics After the Digital Image*. Routledge.