

Integrating Generative AI into Landscape Architecture Education: Methodologies, Applications, and Ethical Considerations

Xun Liu¹

¹University of Southern California, California/USA · xliu8081@usc.edu

Abstract: This study explores the integration of Generative AI into a Landscape Architecture Media course, focusing on the pedagogical strategies, student engagements, and ethical implications. Incorporating tools like Midjourney, Stable Diffusion (with ControlNet and LoRA), and Photoshop's Generative Fill, the course enables students to transcend traditional design paradigms, embracing a future-oriented approach to landscape architecture. This paper discusses the structured module designed to demystify AI concepts, providing hands-on experience, and examines student projects that showcase AI applications in visualizing landscape changes and generating high-quality renderings.

Keywords: Generative AI, design pedagogy, Stable Diffusion, artificial intelligence

1 Introduction

Recent advancements in artificial intelligence (AI) have revolutionized various fields with innovative tools and methodologies, marking a significant shift in creative disciplines, including landscape architecture. Pioneering AI-driven tools such as Midjourney, Stable Diffusion, and the Generative Fill feature in Adobe Photoshop have opened up new horizons for generating a diverse range of visual outputs. These tools, beyond their capacity to generate 2D assets for renderings as outlined by FERNBERG (2023), have been significantly enhanced by the advent of open-source advancements such as ControlNet and LoRA for Stable Diffusion. These advancements empower designers with precise control over image-to-image generation processes. Applications of these tools are diverse and include tasks such as recoloring linework drawings, transforming sketches into detailed renderings, converting 3D models to renderings, and customizing visual styles, among others.

This wave of innovation builds upon earlier foundational work in the field, as evidenced by projects like SWA's "Plaza Life Revisited" (SCHLICKMAN 2020) and extensive reviews by scholars such as CANTRELL et al. (2021) and FERNBERG and CHAMBERLAIN (2023), who have showcased the broad spectrum of AI applications in landscape architecture, from the analytical applications of machine learning to the creative generation of design schemes.

The integration of AI into the curriculum for landscape architecture students represents a nascent yet rapidly evolving field. Given the pioneering nature of this endeavor, the insights presented herein are drawn from firsthand experiences in navigating the challenges and opportunities that accompany the introduction of AI technologies in educational settings. Initially, AI applications were introduced to students through lectures without practical exercises, mainly because the technical requirements for AI – often involving coding or scripting – posed a significant barrier. However, student interest in AI's potential was unmistakably high. By 2023, the availability of user-friendly AI tools, especially for image generation, began to change this dynamic. These tools now produce high-quality images fit for professional use and are significantly more accessible, lowering the entry barriers for students. This

development has opened up new opportunities for integrating AI into our curriculum, enabling hands-on exploration and application.

2 Course Context and Structure

2.1 Overview of the Course

This course, integral to the advanced media curriculum for upper-year Master of Landscape Architecture students in the University of Southern California, diverges significantly from the traditional focus of landscape media courses on representation and digitization techniques. As the realms of design and architecture undergo rapid evolution, there arises a need for a curriculum that not only keeps pace with but also anticipates technological advancements. Addressing this need, the course introduces students to dynamic systems via computational and algorithmic thinking, marking a departure from the generative design methods and parametric tools historically emphasized, such as Rhino and Grasshopper. The strategic inclusion of Generative AI in the 2023 curriculum – alongside other key modules like Patterns and Landforms, Dynamic Systems, and Data-driven Methodologies – acknowledges AI's significant impact and growing relevance in landscape architecture, aiming to furnish students with the requisite skills and knowledge for navigating the changing landscape.

2.2 Structure of the 3-week AI Module

The 3-week AI module of the course was structured into comprehensive three-hour classes that combined lectures, workshops, and discussions. The curriculum was designed to provide practical, hands-on AI experiences without the need for complex coding skills, enabling easy integration into everyday student workflows. Selected Generative AI tools were specifically chosen for their relevance to landscape architecture, forming the course's practical backbone and allowing students to engage with innovative technologies that shape modern design. All tutorials related to the course are authored by the instructor and made publicly available at landscapearchitecture.ai, ensuring wide accessibility and support for students.

Week 1:

Lecture: Introduction to AI

- Overview: This lecture delves into the history, fundamental concepts, and terminologies of AI, along with a comprehensive overview of available tools.
- Objective: To demystify AI for students by simplifying its concepts, thereby making it more accessible and removing any perceived complexity.

Workshop: Midjourney and Adobe Photoshop Generative Fill

- Midjourney: Provides an exploratory platform for text-to-image generation, encouraging students to engage with complex prompts and iterative design processes. Notable for its capacity to spark early-stage design ideas, Midjourney facilitates creative exploration, allowing students to conceptualize and refine landscape features and scenarios.
- Adobe Photoshop (2024 Version): Features the Generative Fill tool, which enhances the image editing process by using AI to replicate and integrate the context of surrounding imagery. This allows for sophisticated manipulation and improvement of visual outputs, ideal for tasks such as adding elements, adjusting textures, or cohesively expanding landscape images.

Week 2:

Lecture: Generative AI in Design

- Overview: This session explores the progression of Generative AI technologies, from Generative Adversarial Networks (GANs) to Diffusion models, and their implementation in image generation across various design disciplines. A comprehensive introduction to the diverse applications of AI in landscape architecture and beyond is provided, showcasing analytical projects, artistic endeavors, and practical professional applications.
- Objective: To inspire students with the expansive range of image generation tools, introducing specific terminologies within the image generation domain and demonstrating how image regeneration serves not only representation and visualization but also research.

Workshop: Stable Diffusion with ControlNet

- Stable Diffusion, a leading text-to-image AI model, is known for producing intricate images based on textual prompts. In combination with ControlNet, this workshop offers students enhanced control over image generation, aligning outputs more closely with their design visions. ControlNet's capacity to set additional parameters on the diffusion process transforms the creative potential, enabling highly precise and intention-aligned visual outputs.

Week 3:

Lecture: AI in Landscape Architecture

- Focus: This lecture delves into the specialized applications of AI within landscape architecture, spotlighting academic research and broader implementations beyond simple image generation. Topics include AI's role in garden/park layout generation, planting design, and topographic modeling, with a focus on understanding landscape dynamics through AI.

Workshop: Training Customized LoRA Model

- LoRA stands as an advanced method for refining AI models, permitting students to adjust Stable Diffusion outputs for landscape architecture-specific needs. This session guides students through the customization process, from parameter adjustment to emphasize desired styles and elements to achieving tailored visual representations for their projects.

3 Students' Work and Reflection

This section covers the weekly assignments for this 3-credit course designed with a six-hour workload per exercise. Throughout the course, student engagement with AI tools varied, reflecting a journey from initial apprehension to proficient use in landscape architecture. The combination of hands-on and theoretical learning enabled students to actively explore AI technologies, with many showing enthusiasm for the diverse capabilities of the tools in design and visualization. Students experimented with AI across various exercises.

3.1 Exercise 1: Exploring Generative AI and Image Editing

In Exercise 1, students engaged with Midjourney to create visual narratives, unleashing AI's potential for generating diverse landscape scenarios. Meanwhile, Photoshop's Generative Fill served as a tool for refining digital renderings, illustrating AI's efficiency in revising design processes. Midjourney was favored for its user-friendly interface, requiring no installation and delivering quick, high-quality results. Yet, it sometimes produced misleading visual outputs, as exemplified in Figure 1, where a student's attempt to depict a "chaparral" garden resulted in images that did not match the specified plant palette.



Fig. 1: Midjourney Exercise (by Kavya Gudihal)




Fig. 2: Photoshop Exercise for iterative approach (by Jordan Fucci)

Photoshop's Generative Fill was recognized for its transformative impact on daily workflows, due to its integration into the suite of professional tools already in regular use. It astounded students with its content-aware capabilities to fine-tune textures, modify elements, and expand scenes. As shown in Figure 2, when tasked to "add a pond with reflection," the tool demonstrated an understanding of the surrounding architecture, seamlessly blending new elements to match the original imagery. However, the technology also presented limitations, sometimes producing anomalous structures or limited modifications, highlighting the need for critical oversight when using AI in landscape architecture.

3.2 Exercise 2: Precision in Image Generation

In this exercise, due to the high technical demands of Stable Diffusion with ControlNet, including the need for robust GPU resources, we established a shared cloud server on Runpod. This ensured that all students had equitable access to this advanced AI tool. Through this exercise, students mastered fine-grained control over the image generation process, learning to tailor outputs for specific landscape architecture design tasks.


Figure 3 illustrates a student's foray into integrating a Rhino-to-Stable Diffusion workflow. The student began with a 3D model in Rhino, exporting rendered views along with color-coded material IDs. These elements were then adeptly manipulated using Stable Diffusion's multi-controlNet features, allowing for intricate control over image details, distinguishing between various natural elements and infrastructural components, and refining edges and shapes with Canny model. The inclusion of a Depth model enriched the visual depth, enhancing the realism of the generated images. This student's work ventured into representing the transient aspects of landscapes (Figure 4), using AI to visualize changes through daily and seasonal cycles, showcasing the robust technical capabilities of the tools and the dynamic essence of landscape design.

While Stable Diffusion proved to be a potent tool for redefining traditional rendering workflows, students encountered challenges due to the high GPU requirements, the complexity of cloud server setup, and the steep learning curve presented by the extensive customizable parameters.


Output Using Stable Diffusion:



Inputs for Multiple ControlNet:




Rendered View from Rhino



ControlNet Input 1: Customized Displayed Mode exported from Rhino, using "Segmentation" model.



ControlNet Input 2: Preprocessed image using "Depth" model from Rhino rendered view.



ControlNet Input 3: Preprocessed image using "Canny" model from Rhino rendered view.

Fig. 3: Rhino-to-Stable Diffusion workflow (by Rongkun Lu)

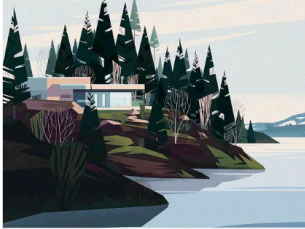







Fig. 4: Representation of Seasonal Changes through adjusting prompts

3.3 Exercise 3: Customizing AI Outputs with LoRA

LoRA sessions gave students hands-on experience in model fine-tuning to customize AI outputs for specific landscape designs. They gathered and labeled 30 images to train their LoRA models, enhancing the AI's ability to generate designs aligned with their visions. Comparative tests were conducted to assess the personalized models' efficacy. Although the focus was on familiarization rather than extensive model refinement, the outcomes, as seen in Figure 5, demonstrated that well-trained customized models could yield high-quality results.

Training Dataset Samples

		
architecture,illustrative drawing, house, trees, plants, vibrant, colorful, landscape,texture, waterside, snow, mountain,cliff, forest	illustrative drawing, house, trees, plants, vibrant, colorful, landscape,texture, waterside, mountain, cliff, forest	architecture,illustrative drawing, house, trees, plants, vibrant, colorful, landscape,texture, waterside, snow, mountain
		
architecture,illustrative drawing, house, trees, plants, vibrant, colorful, landscape,texture, waterside, mountain, cliff, forest	architecture,illustrative drawing, house, trees, plants, vibrant, colorful, landscape,texture	illustrative drawing, house, trees, plants, vibrant, colorful, landscape,texture, waterside, mountain, cliff, forest

...
30 pairs of images and prompts

Using Customized LoRA in Stable Diffusion




<p><lora:illustration:1> CustomizedIllustration, Sunken circular sitting area, couch, fire pit, yard, illustrative style</p>	<p><lora:illustration:1> CustomizedIllustration, Minimalist house in the snow with mountains in the background, illustrative style</p>	<p><lora:illustration:1>CustomizedIllustration, The blue house at the end of the sea, the palm trees, the waves,illustrative style</p>
		

Fig. 5: Customized LoRA model (by Gloria Gao)

4 Conclusion

The course facilitated a comparative evaluation of various AI tools, assessing their utility in achieving educational objectives. Midjourney excelled in enabling students to rapidly conceptualize and creatively explore design possibilities, while Photoshop's Generative Fill was lauded for its seamless integration into existing workflows, streamlining image editing tasks. Stable Diffusion, when combined with ControlNet, delivered a structured method for generating images, aligning closely with specific design needs. Additionally, LoRA's model fine-tuning capabilities opened avenues for personalized applications, allowing for AI outputs to be tailored to the unique demands of landscape architecture projects. Ethical consideration was also important throughout:

Intellectual Property and Authorship: The course catalyzed discussions on how to properly attribute AI-generated content. While there is no one-size-fits-all answer, students were encouraged to diligently cite all inputs, including images, prompts, tools, algorithms, and any community-sourced models. Efforts to use AI-generated images to train customized models were also explored as a means to mitigate credit issues, though these practices remain part of an ongoing dialogue.

Accuracy and Misinformation: AI's propensity to generate incorrect information, such as inaccurate plant species or distorted structural elements, necessitates a rigorous verification process. These errors, while challenging, also served as unintended prompts for students to consider alternative designs and solutions, fostering a space for innovative thinking.

Labor and Automation: While students were excited about AI's capabilities, there was concern about its potential to displace jobs. Reflecting on the adoption of technologies like AutoCAD, it was agreed that AI would enhance but not replace the essential creative and analytical skills in landscape architecture.

By addressing these ethical considerations, the course aimed to prepare students not only technically but also to engage with the broader implications of AI in their professional futures, ensuring a responsible and informed use of technology.

References

- CANTRELL, B., ZHANG, Z. & LIU, X. (2021), Artificial Intelligence and Machine Learning in Landscape Architecture. In: *The Routledge Companion to Artificial Intelligence in Architecture*. Routledge.
- FERNBERG, P. & CHAMBERLAIN, B. (2023), Artificial Intelligence in Landscape Architecture: A Literature Review. *Landscape Journal*, 42 (1), 13-35.
DOI: <https://doi.org/10.3368/lj.42.1.13>.
- FERNBERG, P., GEORGE, B. H. & CHAMBERLAIN, B. (2023), Producing 2D Asset Libraries with AI-powered Image Generators. *Journal of Digital Landscape Architecture*, 8-2023.
<https://doi.org/10.14627/537740020>.
- SCHLICKMAN, E. (2020), Assessing Automation: Methodological Insights from Experimenting with Computer Vision for Public Life Research. *Journal of Landscape Architecture*, 15 (3), 48-59. DOI: <https://doi.org/10.1080/18626033.2020.1886515>.

