Customizable AI Tools for Landscape Architecture: Tailored Processes Compared with Traditional Rendering

Afshin Ashari[1], Ramtin Shafaghati[2]

[1]University of Guelph, Ontario/Canada · aashari@uoguelph.ca
[2]University of Guelph, Ontario/Canada

Abstract: Artificial intelligence (AI) is transforming landscape architecture by enhancing visualization processes and streamlining workflows. This study explores the potential of AI-driven tools, particularly text-to-image diffusion models such as Stable Diffusion and platforms such as ControlNet and LookXAI, in comparison with that of traditional human rendering techniques. Using "The Meadow" an award-winning landscape project, as a case study, AI-generated versus human-rendered visuals are evaluated on the basis of criteria such as geometric accuracy, material fidelity, and planting realism. Quantitative metrics, including cosine similarity and the structural similarity index (SSIM), are employed to assess the output quality.

The results indicate that while AI tools show promise in replicating complex spatial geometries and plant diversity, limitations such as material inconsistencies and biased outputs persist. Ethical considerations, including transparency, data bias, and accessibility, are discussed to ensure the equitable adoption of AI tools in landscape architecture. The study concludes with recommendations for integrating AI methods into design workflows, emphasizing the need for hybrid approaches that balance efficiency with artistic control. By advancing AI applications in the field, this study aims to optimize design processes and improve client collaboration.

Keywords: AI visualization, text-to-image, design tools, ControlNet, landscape architecture

1 Introduction

Artificial intelligence (AI) is transforming the field of landscape architecture by introducing tools that enhance visualization and simplify workflows. These tools enable designers to efficiently create realistic renderings and test various design scenarios, significantly advancing traditional practices (FLORIAN et al. 2024). For example, text-to-image models such as Stable Diffusion are designed to generate images on the basis of textual descriptions. These models refine their outputs through multiple stages, progressively matching the generated visuals with the input description. This iterative process makes it easier for designers to explore creative ideas quickly and effectively (ZHAO et al. 2024, HO et al. 2020).

Platforms such as ControlNet enhance this process by allowing users to guide AI models with reference images or component outlines, ensuring that the generated visuals maintain geometric accuracy and align with specific design requirements. ComfyUI provides a user-friendly interface for managing AI image generation workflows, offering modular tools to adjust various parameters and customize outputs for specific projects. LookXAI, a SketchUp extension, focuses on generating detailed plant textures and materials, making it especially useful for landscape architecture applications where precision in vegetation representation is critical (KARADAG 2023, MA & ZHENG 2024).

Techniques such as prompt-to-prompt editing allow designers to modify AI-generated images by adjusting specific elements in the textual input, offering greater flexibility and control over the final output. Similarly, low-rank adaptation (LoRA) fine-tunes pretrained AI

Journal of Digital Landscape Architecture, 10-2025, pp. 653-664. © Wichmann Verlag, VDE VERLAG GMBH · Berlin · Offenbach. ISBN 978-3-87907-754-0, ISSN 2367-4253, e-ISSN 2511-624X, doi:10.14627/537754062.

models to better suit particular tasks or design aesthetics without requiring extensive computational resources. These methods enhance the precision and customization of AI-generated visuals, making them more relevant to the specific needs of landscape projects.

To evaluate the quality of these outputs, metrics such as the mean squared error (MSE) and structural similarity index (SSIM) are commonly employed. The MSE measures the pixel-by-pixel difference between two images, with lower values indicating greater similarity. The SSIM, on the other hand, evaluates structural and perceptual similarities between images, considering factors such as luminance, contrast, and texture. These metrics ensure that AI-generated images meet the desired standards and align with design objectives (WANG et al. 2004, KIM & YE 2021).

While AI tools offer efficiency and innovation, traditional rendering methods remain crucial for achieving high levels of artistic control and customization. A hybrid approach that integrates AI-generated outputs with manual refinements achieves practical balance, combining the speed and automation of AI with the creativity and expertise of human designers (LI & AMOROSO 2023). This observation was applied throughout various stages of the workflow, ensuring that AI capabilities complement, rather than replace, human input.

This study explores the application of AI tools in landscape architecture, focusing on an award-winning project to compare AI-generated visuals with traditional human-rendered images. By evaluating the strengths and limitations of these tools, this study offers insights into their practical utility and addresses ethical considerations such as transparency, bias, and accessibility to guide the responsible adoption of AI in the field.

2 Methods, Tools, and Techniques

2.1 Research Workflow

This study followed a structured approach to evaluate how closely AI-generated visualizations could resemble human renderings (Fig. 1). The process began by selecting an ASLA award-winning landscape architecture project. Since the goal was to determine how accurately AI could replicate human visualization, a high-quality human-rendered image was created to serve as a reference. This rendering needed to closely match the original project's appearance.

For AI visualization, textual descriptions (prompts) and reference images were needed. Crafting effective prompts involves a process similar to that in reverse engineering – analyzing human-rendered images to extract descriptive keywords that guide AI image generation. This step was crucial for ensuring that AI-generated images closely matched human visualizations. To increase the accuracy of the AI outputs, multiple platforms and tools were tested, refining the generated images for the highest resemblance to human renderings.

A quantitative comparison was necessary to measure similarity, as subjective human evaluation alone was insufficient. To achieve this goal, machine learning and AI-based algorithms were employed to compare images numerically. Throughout this process, prompt engineering and procedural fine-tuning were iteratively refined until the greatest similarity between AI-generated and human-rendered images was achieved.

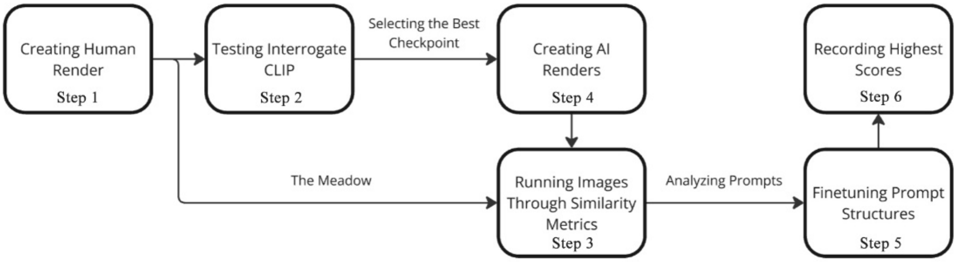

Fig. 1: Workflow diagram illustrating the research methodology for comparing AI-generated and human-rendered landscape architectural visualizations

2.2 Case Study Selection Criteria

To ensure credibility and relevance, we conducted an extensive search through various landscape architecture competitions, ultimately selecting an American Society of Landscape Architects (ASLA) award-winning project. The selection process focused on two key criteria: the project's geometry and its alignment with contemporary design contexts. These factors are critical in evaluating the effectiveness of AI-driven methodologies in real-world applications.

The first criterion, geometry, was particularly important, as AI models often struggle with complex designs owing to their inherent limitations. Preliminary testing across multiple projects revealed that AI tools work more effectively with simpler geometries. For this reason, we chose a project with relatively straightforward design elements. This choice allowed us to refine our methods and establish a solid foundation before applying the approach to more intricate and complex designs.

The second criterion was the relevance of the project to current design trends. It was essential to test the proposed framework on a recently built, contemporary project to assess the effectiveness of AI within a current landscape architecture context. "The Meadow" at the Old Chicago Post Office, designed by Hoerr Schaudt Landscape Architects, met this requirement (Fig. 2). Its manageable complexity facilitated effective AI testing while also providing a clear and relevant reference for comparison with human-rendered images.

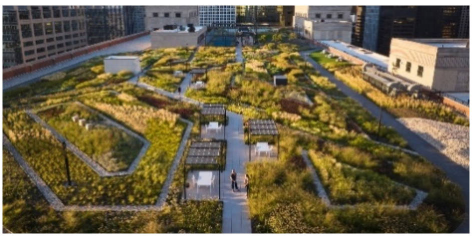

Fig. 2:
"The Meadow" at the Old Chicago Post Office, designed by Hoerr Schaudt Landscape Architects

2.3 Checkpoint Selection and Evaluation

The AI models used in Stable Diffusion operate through "checkpoints" which are specific versions of pretrained models that store learned information about image generation. A checkpoint in AI refers to a saved state of a model at a particular point in its training process, influencing the style, detail, and realism of generated images.

Several checkpoints were tested to determine which model produced the most accurate and realistic landscape visualizations. The RealVisXL checkpoint within ComfyUI was selected for its compatibility with ControlNet and its ability to generate highly detailed vegetation. Checkpoint evaluation was based on the following criteria:

- Geometric accuracy: The ability of the AI tool to replicate spatial relationships and form accurately.
- Material fidelity: The consistency of textures and surface properties across the visualization.
- Planting realism: The accurate representation of plant species and their arrangements.

2.4 Tools and Technologies

Various AI tools and platforms were tested during the initial phases of the study to identify those that provided the best results, allowing for further exploration and deeper understanding. Additionally, this process helped determine which tools are currently not well suited for landscape architecture.

2.4.1 Initial Image References

ControlNet models require structured reference images to guide AI-generated outputs. Three types of reference images were used:

- **3D models from Rhino with Photoshop-applied textures** to ensure geometric structure.
- **Hand or iPad sketches** to guide stylistic interpretations.
- **Real-site photographs** to validate the accuracy of the AI-generated outputs.

2.4.2 Interrogate CLIP

Contrastive language-image pretraining (CLIP) is a deep learning model developed by OpenAI that associates textual descriptions with visual content (BRADE et al. 2023). It is designed to understand the relationship between images and text, allowing for keyword generation from images or image retrieval on the basis of text input.

In this study, CLIP-based tools were tested for generating prompts. ComfyUI's CLIP Interrogator, Midjourney's "/describe" feature, and GPT-4o were assessed for their ability to convert images into text-based descriptions (Step 2 in Figure 1). ComfyUI's CLIP Interrogator produced inaccurate and unrelated descriptions, making it unsuitable for landscape-specific applications. Conversely, Midjourney's "/describe" and GPT-4o provided more precise and contextually relevant descriptions, making them more effective for prompt engineering. Microsoft Bing's Copilot outperformed ComfyUI's CLIP Interrogators in image-to-text accuracy.

2.4.3 Stable Diffusion and SketchUp Extension (LookxAI) (Step 4 in Figure 1)

Stable Diffusion was employed for AI image rendering via ComfyUI, which was chosen for its advanced customization through node adjustments, checkpoint integration, and LoRA. Compared with Stable Diffusion Online and similar tools, ComfyUI offers greater control over rendering, enabling precise modifications for tailored outputs.

LookxAI, a SketchUp extension, was chosen for its ability to generate realistic plant details and textures essential for landscape renderings. Its seamless integration with design workflows and high-quality outputs makes it crucial for assessing the visual similarity of AI-generated images to human renderings.

2.4.4 Final Package and Reason Offer Each Step

The final ComfyUI workflow integrates nodes for tasks such as model loading and prompt input. The process begins with the checkpoint and ControlNet's Canny Edge feature, which is a popular edge detection algorithm that identifies significant boundaries in an image, preserving structural integrity and geometric features for accurate AI-generated outputs (Fig. 3).

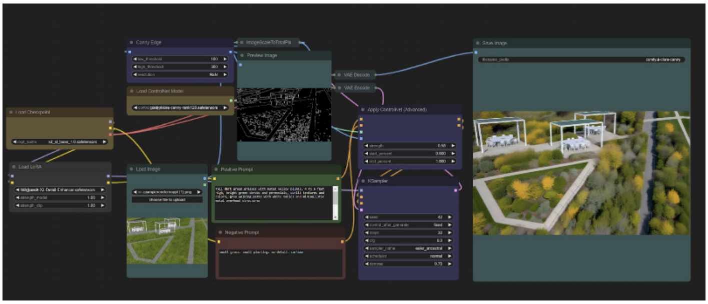

Fig. 3: ComfyUI node configuration displaying the base workflow with checkpoint and ControlNet Canny Edge integration for structural preservation in landscape architectural renderings

An additional workflow incorporates inpainting and upscaling to refine AI renderings. Inpainting fills incomplete areas via targeted prompts, enhancing visual fidelity and alignment with the intended design (Fig. 4). Upscaling increases the resolution for greater detail and clarity, ensuring that the final output meets professional standards for landscape architecture renderings.

(a) (b) (c)

Fig. 4: Advanced ComfyUI workflow incorporating inpainting and upscaling nodes for enhanced detail refinement. (a) Selecting the area for inpainting (b) Inpainting results (c) Upscaling the render to increase detail and clarity.

3 Human Image Rendering

3.1 Selection of The ASLA Winner Project Frame

Given the complexity and hidden layers involved in AI-generated images, a simplified frame with minimal geometric complexity and fewer plant species was initially chosen. This approach allows for an in-depth understanding of AI rendering behavior before expanding to more complex scenes. The primary objective was to establish whether AI-generated visualizations could achieve the same level of realism as that of human-rendered images. By first focusing on a simple structure, we refined the workflow and methodologies before applying them to broader views and additional projects.

3.2 Remodeling in Rhino 3D and Rendering in D5 Render (Step 1 in Figure 1)

(a) (b)

Fig. 5: (a) Human-rendered base model in Rhino 3D, showing the simplified configuration of "The Meadow" project with basic vegetation and architectural elements and (b) Final D5 render output, demonstrating refined lighting, materiality, and advanced vegetation placement for evaluating AI output quality

The Rhino 3D base model (Fig. 5, a), which is based on the constructed project, features basic vegetation, paved areas, tables, and pergolas. This model was imported into D5 Render for refinement. Additional assets, such as shrubs and grasses, were added to match real-site photos. Postprocessing in D5 Render involved adjusting the lighting, hue, and temperature, resulting in a professional, client-ready render (Fig. 5, b).

3.3 Criteria for High-Quality Image Rendering

Since the goal of this study was to transform AI-generated images into human-equivalent visualizations, the evaluation criteria focused solely on human-rendered images. Factors such as plant diversity, material realism, and lighting accuracy were used to assess human render quality. On the AI side, multiple platforms and techniques were tested, with continuous prompt refinement to maximize similarity to the human render.

4 **Comparison and Evaluation** (Step 3 in Figure 1)

4.1 Brief Explanation of Comparison Algorithms

Four algorithms were used to evaluate the AI image generators. The mean squared error (MSE) measures pixel intensity differences, with lower values indicating greater similarity (Wang et al. 2004). The structural similarity index (SSIM) measures structural, luminance, and contrast changes and ranges from 0 to 1. Both metrics require images of the same dimensions and minimal noise for accuracy. Cosine similarity, which is effective in machine learning tasks (Xia et al. 2015), and Euclidean distance, which measures spatial differences

(MATHEWS et al. 2022), are more effective for prompt testing. Slight size mismatches can reduce cosine similarity to 89%, and scaling AI images to match humans often results in lowered scores

4.1.1 Similarity Criteria for Algorithms

Similarity criteria include the AI generator's ability to replicate the render's form, element positioning, and plant species size and color. The algorithms evaluate similarity by comparing attributes such as color, texture, and shape. Cosine similarity, which is used in this study, focuses on the orientation of feature vectors and measures content similarity regardless of intensity, making it preferable to other metrics.

4.1.2 Comparative Analysis of AI vs. Human Rendered Images

Prompts were tested in ComfyUI and LookxAI, with resulting images compared to the human-rendered images using similarity algorithms. Through trial and error, key phrases and the word order were refined to improve similarity scores. Cosine similarity and the Euclidean distance were applied to all prompts, whereas all four metrics, including the SSIM and MSE, were applied to the top three highest-scoring prompts.

4.1.3 Success Metrics and Observations

Despite high similarity scores – 79% cosine similarity and 142.10 Euclidean distance – the LookxAI output exhibited issues, such as a distorted pathway on the right that failed to match the correct materials (Fig. 6).

Fig. 6: ComfyUI-generated landscape visualization highlighting both the capabilities and limitations of AI rendering in replicating complex spatial relationships

5 Results

As different prompts were iterated, their similarity scores were recorded, and the combinations of keywords that produced higher results were noted. These prompts were then modified and fine-tuned to achieve even better outcomes.

5.1 Summary of the Findings (Step 5 in Figure 1)

Following the modification of the base prompt in LookXAI, the similarity metrics indicated a noticeable improvement. The cosine similarity, which ranges from 0 to 1, increased from 74% to 79%, indicating a closer alignment between the AI-generated and reference images (higher values represent better similarity). The Euclidean distance, a measure of spatial difference, decreased from 151.35 to 146.04, reflecting a closer match between the images (lower values indicate better similarity). The mean squared error (MSE), with a range from 0 to infinity, improved from 6831.63 to 6316.06, suggesting a reduction in the error between the images (lower values indicate better accuracy). The structural similarity index (SSIM), which ranges from 0 to 1, slightly increased from 0.037 to 0.039, indicating a minor improvement in structural similarity (higher values represent greater similarity).

Table 1: Comparison of Base and Modified Prompts in LookXAI with Corresponding AI-Generated Images

Prompt	AI Generated Image
Base Prompt: Lush meadow overfilled with short miscanthus sinensis grass with plume-like wheat heads, golden ligustrum shrubs, dark green planting, **dense planting, abundant planting,** short miscanthus Sinensis grass, white concrete pathways, white rectangular pergolas and tables, white chairs, detailed planting, ultrarealism, detailed shadows, natural lighting, **realistic render style**	
Modified Prompt: Verdant meadow overfilled with short miscanthus sinensis grass with plume-like wheat heads, golden ligustrum shrubs, dark green planting, **highly concentrated planting,** short miscanthus Sinensis grass, white concrete pathways, white rectangular pergolas and tables, white chairs, detailed planting, ultrarealism, detailed shadows, natural lighting, **unreal engine 5 render**	

After the base prompt in ComfyUI was modified, several similarity metrics improved. The cosine similarity increased from 0.61 to 0.69, whereas the Euclidean distance decreased from 202.24 to 182.71, indicating a closer match. The MSE improved from 6786.99 to 6029.95, and the SSIM slightly increased from 0.044 to 0.047. These changes suggest that adjustments to the prompt led to a more accurate alignment between the AI-generated and reference images.

Table 2: Comparison of Base and Modified Prompts in ComfyUI with Corresponding AI-Generated Images

Prompt	AI Generated Image
Base Prompt: Verdant meadow overfilled with native perennials, and shrubs, tall grass, dark green planting, white concrete pathways, white chairs, white rectangular pergolas, white roof, detailed planting, hyperrealism, hyperrealistic shadows, natural lighting, unreal engine 5 render	
Modified Prompt: Verdant **light green** meadow overfilled with native perennials, and shrubs, tall grass, **light green** planting, white concrete pathways, white chairs, white rectangular pergolas, white roof, **scattered miscanthus sinensis grass with yellow plume-like wheat heads,** detailed planting, hyperrealism, hyperrealistic shadows, natural lighting, unreal engine 5 render, **landscape photography shot by Sony a7 IVA camera**	

5.2 Evaluation of AI Tools' Effectiveness (Step 6 in Figure 1)

The AI tools achieved satisfactory results, with cosine similarity scores of 0.797 for LookXAI and 0.696 for ComfyUI. Upscaling and LoRAs, such as JJ's landscape design LoRA checkpoint, improved scores, especially for plantings similar to those in the human render. The scores remained below 80% due to differences in plantings, but replacing them in Photoshop with those from the human render increased the score to 0.821 (Fig. 7).

Fig. 7: Improved alignment achieved by replacing AI-rendered plantings with manual edits (Left: Original AI image, Right: Replaced planting)

5.3 LookXAI vs. ComfyUI

LookXAI produced higher-quality images with detailed planting, achieving a cosine similarity of 0.797 compared with that of 0.696 for ComfyUI. The performance depends on the training data, and LookXAI's Imagen2 model outperforms ComfyUI's best landscape-oriented checkpoint. This gap may narrow as new checkpoints are developed by the community or incorporated into stable AI models.

5.4 AI Rendering Challenges

Visual imagery can highlight key limitations of AI rendering in landscape architecture. Without structured workflows such as ControlNet, AI tools struggle to maintain geometric accuracy, and poor-quality references lead to errors in textures and materials. Even with proper inputs, AI tools often fail to capture details such as plant types, object geometry, lighting, and figure placement, and their outputs lack the stylistic coherence of human renderings (Fig. 8). Side-by-side comparisons can emphasize the need for workflows and high-quality inputs to address these issues.

(a) (b) (c)

Fig. 8: Comparative analysis highlighting the limitations of AI rendering in maintaining geometric accuracy and material consistency in landscape architectural visualization. (a) Inability to maintain geometry, (b) rendering style not aligned with that of the human render, and (c) inability to accurately depict plant types.

5.5 Insights on AI's Capability in Landscape Architecture

The role of AI in landscape architecture is growing, but the field lacks major specific platforms such as Midjourney or DALLE, which are tailored to its needs. However, checkpoints and LoRAs from repositories such as Civit.ai, a platform offering AI tools and models for creative applications, and Hugging Face, known for its extensive collection of machine learning models, provide accessible landscape-oriented designs. Tools such as D5 Render add AI features, including 3D asset generation and sky/light matching, whereas Adobe offers AI tools for image generation and enhancement. Although specialized AI tools for landscape architecture are limited, increased adoption is likely to drive the development of targeted solutions.

5.6 Potential Benefits and Challenges

AI rendering streamlines high-quality visual production, saves time, enables rapid design iterations, and enhances client communication. Challenges include limited control over details and inconsistent materials or scales, requiring prompt testing for uniformity.

The proposed workflow using ComfyUI and LookxAI enhances design capabilities but adds complexity, such as cognitive load, compatibility issues, and steep learning curves. Streamlined, integrated tools could simplify processes and reduce inefficiencies.

As the field evolves, all-in-one platforms with seamless functionality and enhanced application programming interfaces (APIs), which allow different software systems to communicate and integrate, can improve efficiency. Collaboration between developers and designers should focus on creating powerful, user-friendly tools tailored to landscape architecture.

5.7 Ethical Concerns, Bias, Accessibility and Equity

AI-generated imagery raises ethical concerns related to bias, accessibility, and transparency. Bias in AI models stems from the datasets used for training, which may lack diversity in landscape architecture representations. This can lead to inaccuracies in generating site-specific elements or culturally significant design features. To mitigate bias, this study emphasizes dataset diversification and the periodic auditing of AI-generated outputs to ensure fair and precise representation.

Accessibility is another critical factor. Advanced AI visualization tools often require substantial computing power and technical expertise, limiting their availability to professionals with access to high-performance hardware. Expanding open-source alternatives and user-friendly interfaces can bridge this gap, promoting equitable AI adoption within landscape architecture.

Transparency in AI-generated visualizations is vital for trust in design practices. Clearly, documenting the tools, models, and processes used in AI image generation allows for reproducibility and accountability. In this study, transparency was maintained through the detailed reporting of checkpoints, prompt engineering methods, and algorithmic comparisons. Establishing standard guidelines for AI-assisted design processes will further enhance clarity and ethical responsibility in landscape architecture visualization.

Another important ethical consideration is the proper crediting of AI-generated images. As AI tools such as checkpoints, LoRAs, and algorithms have become central to landscape design, it is essential to acknowledge the teams and organizations behind these technologies. This includes citing the specific tools and datasets used in image creation. Traditional design practices require the acknowledgment of collaborators. Thus, AI-driven design should follow similar protocols. This approach ensures transparency, promotes ethical AI use, and fosters respect for intellectual property, supporting continued innovation in the field.

5.8 Areas of Future Research and Limitations

Future studies should focus on refining similarity algorithms to align AI-generated images more closely with human perception. A detailed exploration of prompt fine-tuning and keyword adjustments to enhance image similarity to human renders would also be valuable. Given the rapid advancements in AI, current platforms producing high-quality images may soon be surpassed, requiring ongoing evaluation of emerging models.

Additionally, more research is needed on transparency, ethics, and accessibility in AI-generated design. This includes addressing biases in AI models, improving accessibility to AI tools for a broader audience, and ensuring ethical practices in AI-driven workflows.
Compared with traditional methods, incorporating client perspectives is crucial for understanding the acceptance of AI-generated renders. Surveys and interviews could provide rich insights into market expectations, informing both design practices and tool development.

6 Conclusion

This study explored the potential of AI tools in generating landscape architectural visualizations comparable to human renderings. By systematically testing various AI platforms, fine-tuning prompts, and integrating structured workflows, the study demonstrated that AI-generated images can achieve high levels of similarity to traditional human renderings.

AI rendering has the potential to improve design solutions by increasing efficiency, enabling rapid iterations, and increasing accessibility in visualization processes. While AI-generated imagery accelerates early-stage design exploration, human expertise remains indispensable for refining artistic intent and contextual accuracy. A hybrid approach that integrates AI automation with human creativity can optimize workflows and expand the possibilities of landscape architecture visualization.

Further studies should focus on refining AI-generated rendering techniques, improving bias detection, and developing industry-wide frameworks for AI-assisted visualization. As AI continues to evolve, its integration into landscape architecture will require ongoing dialog between designers, researchers, and technologists to ensure its responsible and effective application.

References

BRADE, S., WANG, B., SOUSA, M., SAGEEV OORE & GROSSMAN, T. (2023), Promptify: Text-To-Image Generation Through Interactive Prompt Exploration With Large Language Models. doi:10.1145/3586183.3606725.

DONG, W., XUE, S., DUAN, X. & HAN, S. (2023), Prompt Tuning Inversion for Text-Driven Image Editing Using Diffusion Models. doi:10.48550/arxiv.2305.04441.

GAL, R., ALALUF, Y., ATZMON, Y., PATASHNIK, O., BERMANO, A. H., CHECHIK, G. & COHEN-OR, D. (2022), An Image Is Worth One Word: Personalizing Text-To-Image Generation Using Textual Inversion. doi:10.48550/arxiv.2208.01618.

HERTZ, A., MOKADY, R., TENENBAUM, J., ABERMAN, K., PRITCH, Y. & COHEN-OR, D. (2022), Prompt-To-Prompt Image Editing with Cross-Attention Control. doi:10.48550/arxiv.2208.01626.

HO, J., JAIN, A. & ABBEEL, P. (2020), Denoising Diffusion Probabilistic Models. Advances in Neural Information Processing Systems, 33, 6840-6851. doi:10.48550/arxiv.2006.11239.

KARADAĞ, İ. (2023), Transforming Sketches Into Realistic Images: Leveraging Machine Learning and Image Processing for Enhanced Architectural Visualization. doi: 10.16984/saufenbilder.1319166.

KIM, G., KWON, T. & YE, J. C. (2022), DiffusionCLIP: Text-Guided Diffusion Models for Robust Image Manipulation. doi:10.48550/arxiv.2110.02711.

LI, M. & AMOROSO, N. (2023), An Early Look at Applications for Artificial Intelligence Visualization Software in Landscape Architecture. GIS Open Paper. https://hdl.handle.net/10214/27565.

MA, H. & ZHENG, H. (2024), Text Semantics to Image Generation: A Method of Building Facades Design Based on Stable Diffusion Model. 24-34. doi:10.1007/978-981-99-8405-3_3.

MATHEWS, A., SEJAL, N. & VENUGOPAL, K. R. (2022), Analysis of Content-Based Image Retrieval Using Deep Feature Extraction and Similarity Matching. International Journal of Advanced Computer Science and Applications, 13 (12). doi:10.14569/ijacsa.2022.0131277.

RUIZ, N., LI, Y., JAMPANI, V., PRITCH, Y., RUBINSTEIN, M. & ABERMAN, K. (2023), Dreambooth: Fine-Tuning Text-To-Image Diffusion Models for Subject-Driven Generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 22500-22510. doi:10.48550/ARXIV.2208.12242.

WANG, Z., BOVIK, A. C., SHEIKH, H. R. & SIMONCELLI, E. P. (2004), Image Quality Assessment: From Error Visibility to Structural Similarity. IEEE Transactions on Image Processing, 13 (4), 600-612. doi:10.1109/tip.2003.819861.

XIA, P., ZHANG, L. & LI, F. (2015), Learning Similarity with Cosine Similarity Ensemble. Information Sciences, 307, 39-52. doi:10.1016/j.ins.2015.02.024.

ZHAO, L., SONG, D., CHEN, W. & KANG, Q. (2024), Coloring and Fusing Architectural Sketches by Combining a Y-Shaped Generative Adversarial Network and a Denoising Diffusion Implicit Model. Computer Aided Civil and Infrastructure Engineering, 7, 14678667. doi:10.1111/mice.13116.

ZWANGSLEITNER, F., HABJANIC, G. & KNEGENDORF, A. (2024), AI as a Tool in the Landscape Architecture Design Process. Journal of Digital Landscape Architecture, 9-2024, 987-994. doi:10.14627/537752093.