The Landscape Architect as Developer and Designer: Optimisation and Control of Diffusion Models for Landscape Architecture

Franziska Mees[1], Olaf Schroth[2]

[1]Eastern Switzerland University of Applied Sciences (OST)/Switzerland · franziska.mees@ost.ch
[2]Weihenstephan-Triesdorf University of Applied Sciences/Germany · olaf.schroth@hswt.de

Abstract: The integration of Generative Artificial Intelligence (GenAI), particularly image-generating AI, into landscape architecture marks a decisive step in the digital development of this discipline. This work examines how diffusion models can be demonstrably optimised and controlled to support creative and site-specific design processes within landscape architectural workflows. Focusing on the training and fine-tuning of AI models, the study investigates how customised data and proprietary parameters influence the generation of spatially specific images and design results. A conceptual and technical framework developed in this work, adapted to the open-source generative AI model Stable Diffusion, provides the basis for the practical implementation of training using Low Rank Adaptation (LoRA) and DreamBooth as fine-tuning training methods. This study adopts the classical design process as a holistic framework: Three specific learning tasks – location, concept control and look – define the experimental setup and function as identifiable integration points within the design workflow. The AI models are evaluated both qualitatively and quantitatively through these learning tasks. This paper contributes a structured framework for controlling diffusion models in landscape architectural design. It further demonstrates how these control dimensions can be operationalised within practical design tasks. By examining the interaction between human intent and algorithmic interpretation, this work reveals new possibilities for creative control and optimisation within AI-driven design workflows and positions landscape architects as both designers and active developers of their tools.

Keywords: Generative Artificial Intelligence (AI), diffusion Models, landscape architecture design, fine-tuning and model control, low-rank Adaptation LoRA

1 Introduction

This paper examines the potential for generative AI to be adapted for landscape architecture. The aim is to investigate how landscape architects, as interdisciplinary practitioners, can influence the development of AI in their field. The objective of the study is to benefit from customised models in landscape architecture in order of optimise processes and gain better control over outcomes. The overarching research question is: “How can AI models be customised and used in collaboration with human designers to accurately capture, manage, and creatively develop site-specific, conceptual, and visual aspects in landscape architecture?”

The present study is distinguished by its specific focus on training and fine-tuning of diffusion models in landscape architecture, emphasising on controllability and optimization. By systematically examining learning tasks, this study demonstrates how AI can be integrated into the landscape architectural context – highlighting concrete points where landscape architects can appropriate AI as a creative tool and exercise greater control over design outcomes. It contributes to research by providing a detailed framework for training, evaluating, and optimizing AI models customised to specific tasks in landscape architecture.

Journal of Digital Landscape Architecture, 11-2026, pp. 75-86. © Wichmann Verlag, VDE VERLAG GMBH · Berlin · Offenbach. ISBN 978-3-87907-770-0, ISSN 2367-4253, e-ISSN 2511-624X, doi:10.14627/537770007.

2 Generative AI-Models in Landscape Architecture

2.1 Custom Models: Optimization and Controllability

In landscape architecture, generative AI methods further facilitate the transition of the traditional process of pen and paper creation of design variants into the digital realm (WALLIS & RAHMANN 2016). AI research is now broad in scope, ranging from the use of image-generating AI for visualisations (LI & AMOROSO 2023), the creation of 2D assets (FERNBERG et al. 2023) and applications of GPT models for plant selection (GEORGEet al. 2024) to the training of AI (GANs and diffusion models) with a focus on planting plans (SENEM & TUNCAY 2024). The aforementioned publications show a clear shift in the focus of research within landscape architecture: studies now go beyond experiments with artificial intelligence and are increasingly concerned with systematic evaluations of its capabilities. Research is thus increasingly focused on the optimization and control of AI models.

The analysis of bias in models shows how important it is to critically examine training processes and databases. Global diffusion models, usually trained on billions of images, produce rather generic images that lack the characteristics of a specific location or specific style (KUDLESS 2023). With the help of latent diffusion models developed by ROMBACH et al. (2022) for example, users can optimise AI-models themselves by incorporating their own curated data. Today, up-to-date finetuning methods such as LoRA and DreamBooth, developed by Google researchers (RUOZ et al. 2022), allow to finetune global models such as Stable Diffusion, which has been trained on approximately 5 billion images, to a custom purpose and with a significantly lower number of images.

Studies by HUANG & LI (2024) and the design study on bridge aesthetics by ZHANG et al. (2024) demonstrate that models can be enhanced through the use of proprietary image data and text files or by integrating specific knowledge, such as expert knowledge from the fields of ecology and landscape systems as well as design knowledge. COOP Himmelb(l)au claim in their Instagram account that the “Deep Himmelblau Diffusion Node generative AI Model for Architectural Design outperforms its counterparts Midjourney, Stable Diffusion, and Dalle2 in text2image generation by outputting architectural semantics aligned with Coop Himmelb(l)au’s distinctive formal language” (COOP HIMMELB(L)AU 2025).

Although generative AI models have been widely optimised and controlled, only a few approaches exist in landscape architecture using image-generating AI, highlighting the need for a domain-specific methodological framework to adapt and control AI models in this context.

3 Methods

The research design systematically investigates the training of diffusion models in the context of landscape architecture, with a focus on incorporating the perspective of landscape architecture using fine-tuning techniques. The core objective is to explore the potential of personalised AI models in the early stages of the design process and, beyond that, to customise and implement these models in a practical manner. To achieve this, three distinct research questions are posed, each of which is investigated through targeted training tests and corresponding model evaluations.

The study explores AI models such as Stable Diffusion XL 1.0 using targeted fine-tuning (DreamBooth-Training / Low Rank Adaptation) to generate location-specific details, design-specific knowledge and characteristic graphic styles. To this end, three different training models were developed and tested: the location model, the concept control model and the look model.

Table 1: Research questions (I-III) to define the objectives for targeted model training

Framework	**Research question**	**Resulting model**
I. Site specificity	How accurately can a fine-tuned AI generate site-specific details?	I location model
II. Concept control	How can designers and AI work together to achieve truly novel design outcomes?	II concept control model
III. Look	Can the characteristic visual style (look) of a particular landscape architecture office be learned and regenerated by AI?	III look model

A mixed-methods approach is used, including training Stable Diffusion using pre-experiments and fine-tuning techniques LoRA and DreamBooth. The training process is divided into sequential phases, each reflecting different areas of application relevant to the early stages of landscape architecture (Table 1). The models were trained using specific data sets – adapted text-image pairs – from Heidelberg (I), best practice examples of social space quality and climate adaptation (II), and a specific landscape architecture firm (II/III). The training environment included powerful hardware (RTX 4090 GPU, 64 GB RAM) and software (Kohya_ss GUI/ DreamBooth Extension).

1) Location model (DreamBooth): Clear evaluation criteria are specified for the evaluation of the model to assess the identity-defining characteristics. Twenty images are created for each concept and evaluated on a five-point scale for each typology. A comparative non-parametric analysis (Mann-Whitney U test) is then performed with the same number of images using the base model (SDXL 1.0).

2) Concept control model (DreamBooth): trained on text-image pairs with integrated additional knowledge of places and then interacts with the location model (I) using an additional model merge. Quantitative analysis is limited due to its exploratory nature. Instead, input and output are compared based on specific criteria: image composition, integration of additional knowledge using captions and querying them during image generation (social place quality, infrastructure, climate adaptation). This approach helps to determine whether the model has acquired design knowledge.

   Look model (LoRA): trained on 50 text-image pairs of visualisations. HSL analysis is used in the third training test to quantify the similarity of the generated images to the original style and to gain insights into the efficiency of the training. The images are analysed distinguishing the hue, saturation and lightness (HSL) components to compare their styles. However, other methods such as colour analysis (comparison of RGB values) or texture analysis (examination of patterns and surface details) could also be used in future research. The choice of HSL analysis is therefore not exclusive, but representative of various possible approaches to eval uating and optimising style transfer models.

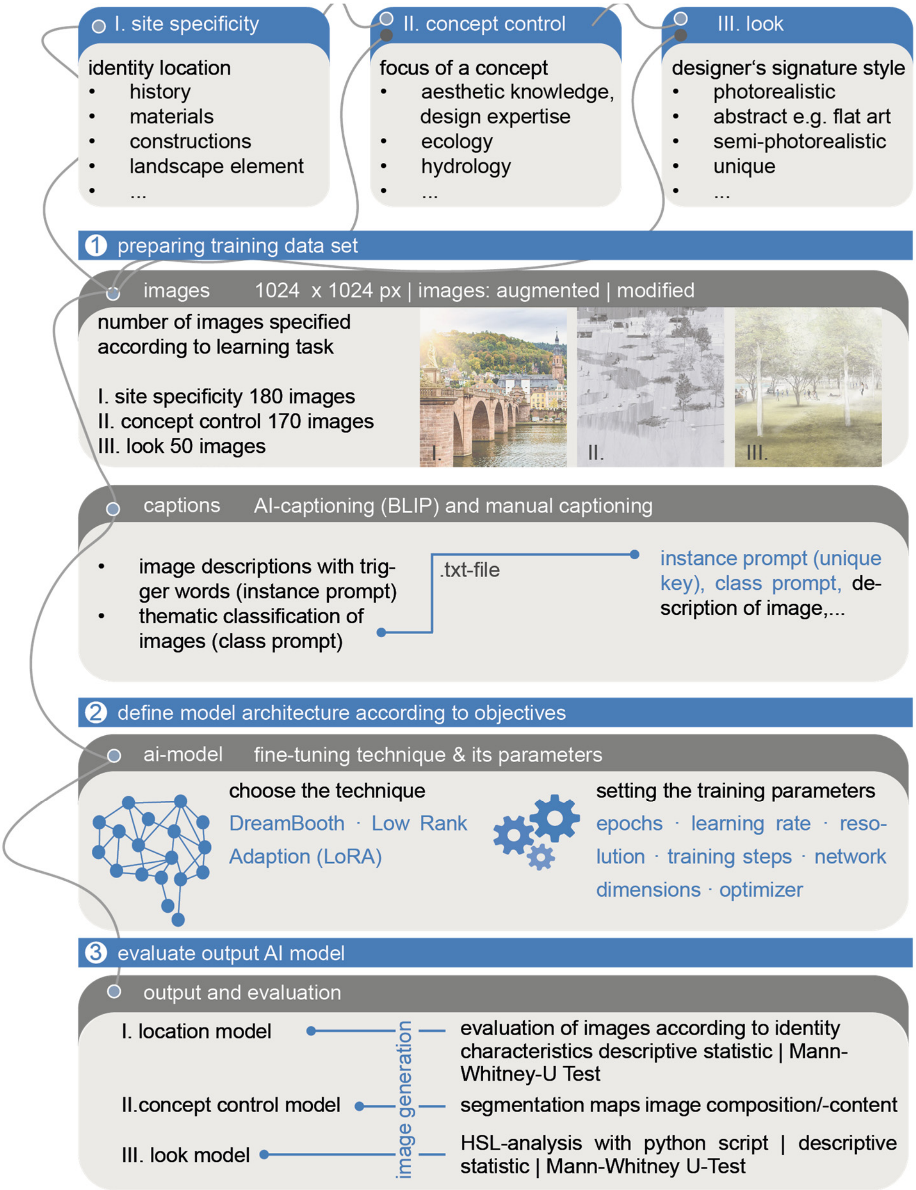

Fig. 1: Methodical training framework for AI training

4 Results

The work demonstrates that generative AI models can learn and regenerate specific expertise, thematic focuses, and graphic styles through targeted fine-tuning. The results highlight the potential of personalised AI for landscape architecture, particularly in the design phase, and open new possibilities for collaboration between humans and AI.

I. Location model: The aim was to create a realistic representation of location-specific details based on a dataset of 180 images from Heidelberg from four identical typologies – architecture (Castle), landscape-integrated elements (Old Bridge), town square (Bismarckplatz), city centre (Old Town). The trained model performed significantly better across three of four typologies (Old Bridge, Castle, and Bismarckplatz) than the base model (Table 2).

Table 2: Quantitative evaluation of the location model

Typology	AI model	Mean	Std. dev.	Median	iqr
Old Bridge	AI trained	2.85	0.93	3	1.75
	Base modell	0.7	0.8	0.5	1
Old Town	AI trained	3.75	0.44	4	0.75
	Base modell	3.4	0.5	3	1
Castle	AI trained	3.45	1.05	4	0.75
	Base modell	1.5	0.61	2	1
Bismarckplatz	AI trained	2.7	0.47	3	1
	Base modell	1.25	0.44	1	0.75

II. Concept control model: The aim was to integrate design-specific knowledge and thematic focal points such as social-place quality, climate adaptation and integrated infrastructure. The model was trained with 170 images and additionally combined with the previously trained location model. The results (Figure 3) show a fusion of location and concept information, which indicate a possible interaction between user and AI.

To analyse how far the fine-tuned model was improved through the curated images from the fine-tuning, input images (left column) were compared to output images (right column) through segmentation and content analysis. Figure 4 is showing the analysis for the prompt "spacious, social place quality, climate adaption".

Fig. 2: Location model: ltr: Old Bridge, Old Town, Castle, Bismarck Square – top: basic model (SDXL 1.0) – row below AI trained

Fig. 3: Result of combining trained models I (location model) and II (concept control model)

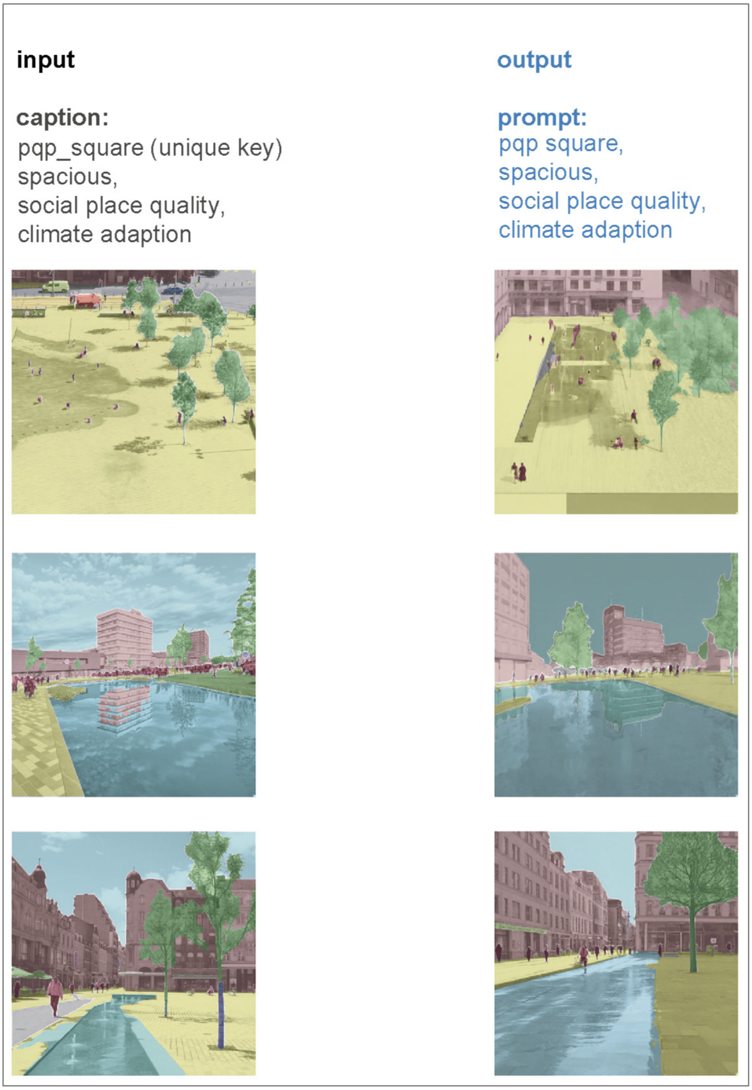

Fig. 4: Concept control model: Qualitative analysis of image content and successful training using segmentation maps

III. Look model: The aim was to learn and reproduce a specific graphic style with a clear focus on efficiency and stylistic accuracy. The model was trained using 50 images, all of which feature a hybrid, sketch-like style, and processed sequentially in several steps. To evaluate the success of the style transfer and training process, an HSL analysis (Fig. 5) was performed. This HSL analysis was carried out using a Python script generated and modified with ChatGPT (OpenAI 2024; ChatGPT (Version GPT-4o) [Large language model]. https://chat.openai.com/). The results showed that after three stages (epoch III), the model came closest to the input style, while a fourth stage (epoch IV) did not bring any significant further improvement.

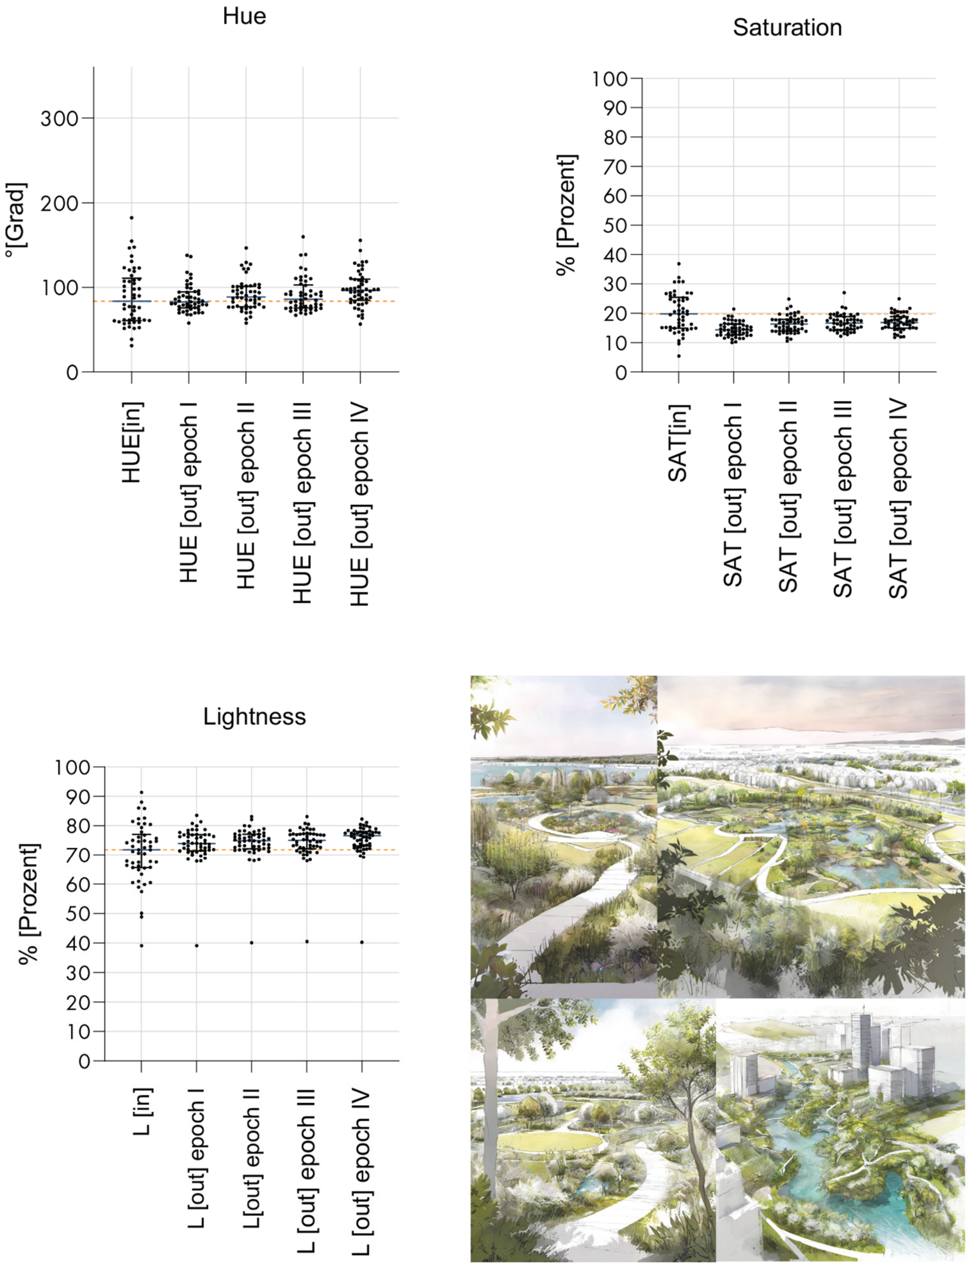

Fig. 5: **L**ook model: HSL analysis, bottom right: results look model (LoRA)

5 Discussion

5.1 Site-Specificity – Precision of Location-specific Generation

How accurately can a fine-tuned AI generate site-specific details?

The results show that finely tuned models can reliably capture location-specific features, although performance depends heavily on data quality and spatial complexity. For example, characteristic open space types, architectural objects and landscape elements in Heidelberg's city centre could be convincingly reproduced. The challenges included capturing complex spatial perspectives and avoiding misinterpretations. Bismarckplatz in particular proved to be a very complex learning task due to its small-scale, multi-layered and at the same time highly distinctive spatial structure, which led to comparatively lower evaluation results. Data preparation was made even more difficult by the limited availability of images and the need to remove distracting elements. In this case, DreamBooth training was used so that the model could learn the entire spatial composition of the location. A prepared training matrix made it possible to teach various spatial concepts in a targeted manner. The model shows clear learning performance: individual objects and perspectives can be addressed directly using specific trigger words, such as characteristic spatial objects. Particularly noteworthy is the training using the example of the castle, where heterogeneous structural design was successfully learned and is reflected in realistic results in image generation.

The work shows that such approaches can be transferred to other forms of location specific knowledge. It is also conceivable to integrate historical structures based on archive images, location-specific textures such as materials, or even self-made sketches from site visits. Such data sets could further support location-based, AI-assisted analysis and design.

5.2 Controlling the Concept – Humans and AI in the Design Process

How can designers and AI work together to achieve truly novel design outcomes?

The results show that guided interaction between human expertise and AI can lead to novel design outcomes. The findings demonstrate that targeted annotation and the combination of different models can integrate individual design knowledge into AI. Conscious control via prompts and training data unlocks creative potential and enables the development of innovative design variants. This approach also offers potential for participatory design processes and integration into education.

Instead of reproducing existing content, the focus was on the generative ability of AI systems to create new design proposals. These results continue to be based on the underlying design concepts and input data. In an additional experiment, concept and location models were combined to explore how different sources of knowledge and design authorship can be brought together. Both AI training and prompting proved to be iterative processes that required important design decisions on the part of the user, including the definition of goals, the degree of influence, and the selection and structuring of training data.

The training focused on spatial representations, especially photographs, as the AI models' prior knowledge in this area is significantly better developed than in abstract forms of representation such as site plans. Captions proved to be essential for embedding expert knowledge into the model. Through fine-tuning, expert knowledge, office-specific preferences and design principles can be specifically utilised using image and text material, thereby achieving a controllable result.

The approach therefore enables designers to generate new insights, ideas and design impulses based on their own text-image data sets. This allows for controlled image generation, the communication of design values and the development of independent design approaches. As a result, the role of the designer is transformed from that of a user to an active co-creator of generative processes.

5.3 Look – Transferability of Stylistic Features

Can the characteristic visual style (look) of a particular landscape architecture office be learned and regenerated by AI?

The results show that stylistic features can be reliably learned and transferred if the training data is consistent and well structured. The Look model has demonstrated that consistent style transfer is possible with stylistically homogeneous training data and carefully optimised parameters. However, excessive deviations or excessive detail in the descriptions can compromise the results. Targeted adaptation of the training data set and the use of suitable model architectures proved to be key factors for successful style adaptation.

The applied HSL analysis provides an initial method for evaluating stylistic consistency but does not fully capture the graphic style. Even a limited number of training cycles proved sufficient to achieve resource-efficient results. Future work should include additional evaluation metrics such as texture features or RGB values and focus on clearly defined stylistic training tasks. Combining a style model (LoRA) with ControlNet proved to be an effective approach for achieving consistent stylistic results in different locations. ControlNet provides contextual control by incorporating spatial information. This ensures spatial coherence while maintaining a distinctive visual style.

6 Conclusion and Recommendations

This study demonstrates that diffusion models can be systematically controlled and optimised along three stages – location, concept control and look – enabling the generation of site-specific and stylistic coherent design outputs aligned with their design goals.

Finely tuned AI models can directly support landscape architects in generating location specific designs while meeting both aesthetic and stylistic requirements in practical workflows. Methodically precise preparation of the training data and active, iterative control of the AI training process by the users are essential for success. Precise goal definition and division into smaller, clearly defined training units are particularly important for complex, context-dependent tasks. The integration of expert knowledge and iterative user participation in the training process enable personalised AI solutions, demonstrating how landscape architects can tailor models to their specific de-sign needs, providing them tangible tools for site-specific-tasks. The qualitatively and quantitatively evaluated results show that AI models can reliably learn and reproduce both location-specific features and self-introduced graphic styles from their own curated datasets. In particular, the ability to incorporate expert knowledge with the help of sample images and technical descriptions, as well as advantages in terms of data protection and competitiveness, open up new perspectives for design and planning in landscape architecture.

As their understanding of the technology grows, landscape architects can train their own customised models to better reflect specific requirements and increase competitiveness. Future research should focus more on how industry-specific knowledge can be integrated into AI models. It also remains to be seen how site plans and other complex formats can be reliably created by AI through targeted training and larger data sets, and how the interaction between image generators and other software should be designed so that designs can be efficiently integrated into workflows.

This study focused on image generators and the underlying diffusion models. Nevertheless, georeferenced work remains a central aspect of landscape architecture practice. Current diffusion models generate raster graphics and rely on comparable data sets. However, the pure diffusion model does not offer a comprehensive solution for location-specific contexts using geodata but requires the integration of additional AI models with different network architectures and the development of suitable interfaces to incorporate and link spatial information in a location-specific manner. The development of an entire workflow and the implementation of diffusion models in established software solutions with spatial reference, such as CAD programmes or GIS systems, may offer another opportunity to develop personalised models and make them accessible to a wider range of users.

Acknowledgements

This article is funded by the Open Access Publication Fund of Weihenstephan-Triesdorf University of Applied Sciences.

Use of LLMs

The following large language models were used only for language improvement purposes: ChatGPT (OpenAI, San Francisco, CA, USA); DeepL (DeepL SE, Cologne, Germany); Microsoft Copilot (Microsoft Corporation, Redmond, WA, USA).

References

COOP HIMMELB(L)AU (2025), Deep Himmelblau Diffusion Node: A Generative AI Model for Architectural Design. https://www.instagram.com/p/CzMM4Vjst5W/ (November 1st, 2025).

FERNBERG, P., GEORGE, B. H. & CHAMBERLAIN, B. (2023), Producing 2D Asset Libraries with AI-powered Image Generators. Journal of Digital Landscape Architecture, 8-2023, 186-194. doi:10.14627/537740020.

GEORGE, B. H., CHAMBERLAIN, B. & FERNBERG, P. (2024), Assessing the Value of Artificial Intelligence in Plant Selection. Journal of Digital Landscape Architecture, 9-2024, 277-284.

HUANG, L & LI, T. (2024), Design Efficacy Evaluation of a Landscape Information Modeling-Stable Diffusion (LIM-SD)-based Approach for Ecological Engineered Landscaping Design: A Case Study of an Urban River Wetland. Landscape Architecture Frontiers, Bd. 12, 68-81, doi:10.15302/J-LAF-1-020103.

KUDLESS, A. (2023), Hierarchies of bias in artificial intelligence architecture: Collective, computational, and cognitive. International Journal of Architectural Computing, 21 (2), 256-279, Juni 2023, doi:10.1177/14780771231170272.

LI, M. & AMOROSO, N. (2023), An Early Look at Applications for Artificial Intelligence Visualization Software in Landscape Architecture. Journal of Digital Landscape Architecture, 8-2023, 543-553, doi:10.14627/537740057.

MEES, F. (2025), Generative KI in der Freiraumplanung. Training von Diffusionsmodellen im Kontext der Landschaftsarchitektur. Bachelorthesis, Weihenstephan-Triesdorf University of Applied Sciences, Freising, Germany.

ROMBACH, R., BLATTMANN, A., LORENZ, D., ESSER, P. & OMMER, B. (2022), High-Resolution Image Synthesis with Latent Diffusion Models, arXiv: arXiv:2112.10752. doi:10.48550/arXiv.2112.10752.

RUIZ, N., LI, Y., JAMPANI, V., PRITCH, Y., RUBINSTEIN, M. & ABERMAN, K. (2022), DreamBooth: Fine tuning text-to-image diffusion models for subject-driven generation. arXiv:2208.12242, doi:10.48550/arXiv.2208.12242.

SENEM, M.O. & TUNCAY, H.E. (2024), Generating Landscape Layouts with GANs and Diffusion Models. Journal of Digital Landscape Architecture 9-2024, 137-144.

WALLIS, J. & RAHMANN, H. (2016), Landscape Architecture and Digital Technologies: Reconceptualising design and making. Routledge, New York.

ZHANG, L, TIAN, X., ZHANG, C. & ZHANG, H. (2024), Aided design of bridge aesthetics based on Stable Diffusion fine-tuning, arXiv preprint arXiv:2409.15812. doi:10.48550/arXiv.2409.15812.