Think AI-side the Box!

Exploring the Usability of Text-to-Image Generators for Architecture Students

Jonathan Dortheimer[1], Gerhard Schubert[2], Agata Dalach [3], Lielle Brenner[4]
Nikolas Martelaro[5]
[1,4] Ariel University, [2,3] Technical University of Munich, [5] Carnegie Mellon University
[1]jonathand@ariel.ac.il [2,3]{schubert|agata.dalach}@tum.de
[4]liellejo.brenner@msmail.ariel.ac.il [5]nikmart@cmu.edu

This study examines how architecture students use generative AI image generating models for architectural design. A workshop was conducted with 25 participants to create designs using three state-of-the-art generative diffusion models and BIM or 3D modeling software. Results showed that the participants found the image-generating models useful for the preliminary design stages but had difficulty when the design advanced because the models did not perform as they expected. Finally, the study shows areas for improvement that merit further research. The paper provides empirical evidence on how generative diffusion models are used in an architectural context and contributes to the field of digital design.

Keywords: *Machine Learning, Diffusion Models, Design Process, Computational Creativity.*

INTRODUCTION

Recent advances in generative machine learning models have demonstrated the potential to revolutionize the field of architecture by introducing a new kind of creative behavior enabled by Artificial Intelligence (AI) tools. For example, Generative Adversarial Networks (GAN) models have already been used in specific architectural tasks such as plan generation (Chaillou, 2020; Zheng et al., 2020), concept image generation (Silvestre et al., 2016; Eroğlu & Gül, 2022) and urban planning (Boim et al., 2022; Zhong et al., 2022). Recently, new image generation diffusion models have demonstrated an extraordinary capability to produce novel and creative architectural illustrations, including architectural designs from text descriptions and images (Ramesh et al., 2021).

It remains unclear how architects can use the diffusion models in architectural practice, which models are more appropriate, and how the use of the models affects the design process. To address this knowledge gap, we conducted an experimental workshop with 25 architecture students to design using three state-of-the-art diffusion models (Dall-E, MidJourney, and Stable Diffusion) and BIM/3D modeling software (Revit). This study's research question is "How could image generating models be used in the context of the architecture studio?" This paper contributes empirical evidence on how architects use generative diffusion models to produce designs and identifies some opportunities and challenges using these models in practice.

AI IN ARCHITECTURE

Architecture constantly evolves due to aesthetic trends and the development of new design tools, such as computer-aided architectural software. These modern tools had a significant impact on the architectural process. Early predictions for computer use in architecture included task automation, alternative work methods, and machine partnership in design evolution (Negroponte, 1972). Drafting software and BIM address the first two visions, while AI tackles the third.

Antoine Picon (2020) wrote: "Artificial intelligence is about to reshape the architectural discipline." In his paper, he pointed out that AI technologies could improve efficiency, help design new things, relieve humans from performing repetitive labor, and support them in creative tasks.

Several approaches address the challenge of AI supporting architects' design processes. While early research, such as Simon (1969), viewed design as a search in the solution space, later studies considered it an exploration (Come et al., 1994). Design problems are often ill-defined, requiring exploration of the problem space (Maher et al., 1996). AI can augment design by extending solution exploration, similar to how multiple human designers approach the same problem (Dortheimer, 2022).

Generative AI

Breakthrough research in machine learning (ML) has demonstrated a high ability to create new images using unsupervised learning (Goodfellow et al., 2014). Different models have produced novel images based on a given text prompt, known as reverse image caption (Mansimov et al., 2016). Diffusion models further improved the text-to-image transition, transforming noise into realistic pictures (Saharia et al., 2022). State-of-the-art models include Stable Diffusion (Rombach et al., 2021), Imagen (Saharia et al., 2022), DALL-E 2 (Ramesh et al., 2021), and MidJourney.

Several approaches have been made to study the use of such technologies in architecture. One of the first approaches was using convolutional neural networks to generate new architectural pictures (Silvestre et al., 2016). Later, DALL-Pytorch was trained to generate new drawings and plans suitable for architecture (Bolojan, Vermisso & Yousif, 2022). In another study, Pix2pixHD was trained using floor plan datasets (Zheng et al., 2020). The researchers were able to generate realistic floor plans in the given style, though not always with functional logic. Other studies involved several GANs sequentially to floor plan generation (Chaillou, 2020).

Another possible application is using text-to-image generators for architectural building images (Zhu et al., 2018; Eroğlu & Gül, 2022). Moreover, GANs may also find their purpose in urban planning. Previous research shows that AI can learn urban patterns and generate resembling patterns (Boim et al., 2022).

However, text-to-image generators require well-crafted text prompts to produce satisfactory results, leading to the development of "prompt engineering" (Oppenlaender, 2022). Still, using text alone for architectural image generation has limitations due to the reductive nature of textual descriptions and the need for explicit architectural training in general-purpose models (Bolojan et al., 2022). Researchers have explored less specific prompts, GPT-3 keyword suggestions (Liu et al., 2022), and generators that incorporate schematic drawings alongside text prompts (Gafni et al., 2022).

Machine and Human Creativity

Creativity could be understood as creating original content (Boden, 2004) but is typically perceived and evaluated subjectively (Mrosla, Koch & von Both, 2019). According to Boden (2004), there are three types of creativity: combinational, exploratory, and transformational. Creativity in artistic forms was already targeted by computational systems, while creative problem-solving received less attention in research (Olteţeanu, 2020).

Several studies explored how humans and AI models co-create design using images as inspiration. While AI-generated images may be perceived as less creative, combining human and machine creativity can enhance innovation (Huang et al., 2021), Research has explored workflow strategies (Yousif & Vermisso, 2022), community collaboration (Epstein et al., 2022), and direct cooperation between designers and AI (Gmeiner et al., 2023).

Other studies claim that AI can benefit low-performing teams (Zhang et al., 2021). The researchers suggested that the communication issues between humans and AI could be solved by referring to human-human cooperation methods (Gmeiner et al., 2023). Nevertheless, it can be assumed that various design problems will require different AI-human collaboration methods and workflows.

METHOD

To answer the study's research question, "How could image-generating models be used in the context of the architecture studio?", the research team hosted the three-day architecture student workshop aimed to expose students to diffusion models and allow them to use the models through an architectural design task to produce a small building using any design tools they choose. The workshop provided a valuable opportunity for the research team to conduct an exploratory study to learn how the students use these models with a considerable design task.

Participants

The workshop was open to architecture students in their third to fifth year of study. A total of 25 students participated, with 18 females and seven males: two were in their fifth year, one was in their fourth year, and the remaining 17 were in their third year. To minimize the pressure on the participants, the workshop was designed to ensure that the student's scores were not dependent on their design performance. Therefore, they were awarded one credit for their participation, whether they produced any design using AI tools or not.

Workshop Structure

In this three-day workshop, participants engaged in a two-phase process that combined lectures and exercises to explore the possibilities, limits, and framework conditions of AI models in architectural design.

Phase 1, the experimental exploration day, was based on Italo Calvino's "The Invisible Cities." Participants selected a city from the book, imagined it, and then visualized it using various AI image generators, such as Stable diffusion, Dall-E, and MidJourney. This introductory exercise aimed to develop the semantics for image generation through a design-based approach. Participants were asked to reflect on their experiences, successes, and challenges while creating the images, comparing their imaginary visions with the AI-generated outcomes, and evaluating the tools and workflows used.

In Phase 2, participants were tasked with employing AI models to generate a three-dimensional building. This phase focused on the autonomous and self-directed implementation of the models for inventive architectural design. The design task involved creating a multicultural community building with two floors, 200 square meters in size, and including stairs to connect the floors. The building supposed to be designed to celebrate the diverse cultures of the local population. This task aimed to make participants deal with the complexities of transforming 2D images into 3D models, incorporating all the necessary systems and components for a functional and culturally inclusive space.

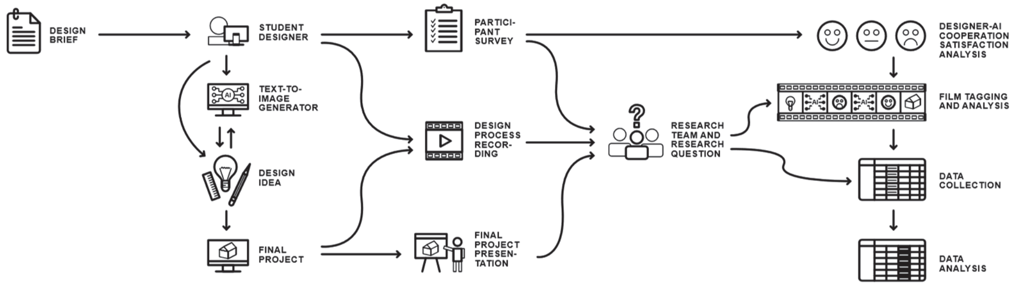

Figure 1
Data collection and analysis

Data Collection

The research data includes final presentation slides and recordings of the participants' computer screens using Zoom software (see Figure 1). The recordings documented the participants' use while working with the AI models and CAAD software. When the participants finished the design, they saved the recording file and provided it to us. According to the Internal Review Board approval, the students were free to participate in any part of the study without affecting their grades.

Data Analysis

The research team viewed, analyzed, and coded the recordings. The coding scheme includes the software used (Dall-E, MidJourney, Stable Diffusion, Revit, and SketchUp), the text prompts the participants provided to the AI models, and the prompt kind (text, image, variation, and image extension). Special notice was provided to the kind of architectural medium the participant requested, such as a plan, section, elevation, or 3D model, and whether the participant used the generated result. When sequential prompts were provided, the team identified the changes between the prompts, the architectural medium requested, the generated output, the changes to the prompt, and if the participant used the output.

RESULTS

The research team concluded the workshop according to plan, having collected 25 student presentations and 16 video recordings with an average duration of 4:02 hours. Three research team members then conducted a qualitative analysis of the video recordings, noting 645 instances of CAD software usage with an average of 40.31 per recording (SD 16.62).

Image Generator Usage and Efficiency

The utilization of AI models was documented in Table 1, with a mean of 36.62 (SD 17.03) instances of usage. MidJourney was the most employed model, with an average of 13.93 (SD 11.35) instances, followed by Dalle-E (M = 13.00, SD = 12.12), Stable Diffusion (M = 6.43, SD = 7.42), and Lexica (M = 2.25, SD = 2.04). Additionally, four participants utilized ChatGPT to facilitate discourse and articulate the text prompts.

The average success rate of AI models was calculated as the average number of successful usage instances divided by the total number of usage instances (see examples in Figures 2, 3, 4). MidJourney had the highest success rate at 39.85% (SD 20.68%), while Dall-E, Stable Diffusion, and Lexica had success rates of 22.88%, 23.93%, and 21.70%, respectively. These results suggest that MidJourney outperformed the other models

Table 1
Summary of model usage instance and computed success rate for each participant

PCPs	Counted Instances				Success Rate			
	Mid Journey	Dall-E	Stable Diffusion	Lexica	Mid Journey	Dall-E	Stable Diffusion	Lexica
1	16	4	5	2	56.25%	50.00%	40.00%	0.00%
2	2	17	23	6	50.00%	52.94%	52.17%	16.67%
3	21	15	1	1	42.86%	40.00%	0.00%	0.00%
4	15	11	3	2	40.00%	9.09%	33.33%	0.00%
5	10	3	7	4	10.00%	0.00%	0.00%	25.00%
7	5	1			0.00%	0.00%		
11	6	3	1	1	83.33%	66.67%	0.00%	0.00%
12	26	31		3	19.23%	12.90%		0.00%
13	23	6	9	1	56.52%	0.00%	22.22%	100.00%
14	5	23	4		40.00%	17.39%	50.00%	
16	19	1	3	2	42.11%	0.00%	33.33%	50.00%
21	17	23		7	47.06%	21.74%		57.14%
22		25	15	3		16.00%	26.67%	0.00%
23			20				45.00%	
26	15	39	12	3	33.33%	23.08%	8.33%	33.33%
27	43	6		1	37.21%	33.33%	0.00%	0.00%
Avg.	**13.94**	**13.00**	**6.44**	**2.25**	**39.85%**	**22.88%**	**23.93%**	**21.70%**
SD	**11.35**	**12.12**	**7.43**	**2.05**	**20.68%**	**21.42%**	**20.21%**	**31.17%**

regarding participants' satisfaction with the outcomes.

Architectural Prompt Anatomy

This study additionally provides an analysis of the terms used to generate architectural images. Specifically, three categories of terms were identified: object properties, situation description, and image properties. Object properties encompass building usage, materials, building details, building size, architectural style, and building form. Examples of object properties include building usage (e.g., community center, art center, entrance lobby), materials (e.g., stone, concrete, glass, iron, colors), building details (e.g., Arabic ornaments, stone carvings, inner courtyard, vertical shutters), building size (e.g., two stories, 200 square meters, eight meters high), architectural style (e.g., modern, architect name, futuristic, crystal palace, Arabic), and building form (e.g., teardrop, organic shape).

Situation description includes building placement (e.g., underground, in the field), building relationship (e.g., facing each other, single building, two buildings, next to a road), architectural scale (e.g., 1:00, 1:500), and geographic location (e.g., Singapore, New York, Melbourne).

Image properties encompass architectural medium kinds and photograph properties. Examples of image properties include architectural medium kinds (e.g., 3D model, plan, section, elevation, sketch, render image) and photograph properties (e.g., wide angle, axonometric, bird view, Lumion render, ultra-HD). The research team observed several intriguing applications of text prompts. One participant employed the phrase "correlated plan and section" to generate 3D renders with plans and sections, which was beneficial. Additionally, another

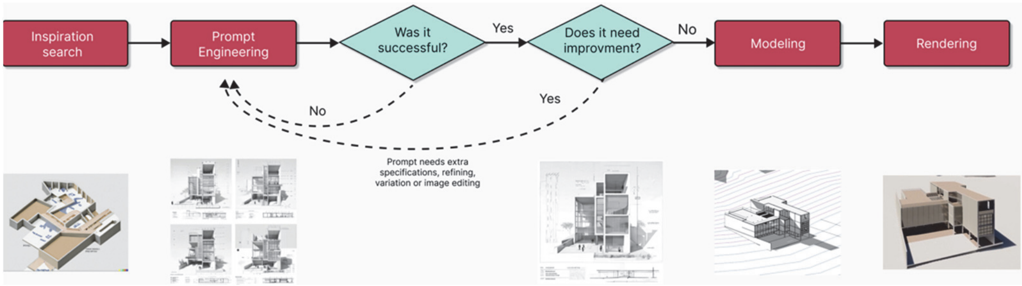

Figure 2
A typical design workflow of one of the participants

participant combined the styles of two architects into a single edifice by utilizing their respective names.

The qualitative analysis of the 25 student presentations revealed a greater understanding of the participants' perspectives on utilizing the models. The presentations also enabled the research team to ascertain which generated images were eventually employed by the participants. This analysis enabled the research team to discern the design activities in which the models were utilized and the efficacy of their implementation (see Figure 2). All the participants employed the AI models to explore design ideas based on the design brief text. The participants were free to select any AI models, with some opting for a single model and others utilizing several models. To further develop the design idea, some participants used the `variations' feature or uploaded images with or without a text prompt to generate new variations.

In the later stages of the design process, some participants attempted to use AI to generate two-dimensional architectural drawings of plans, sections, and elevations. However, these efforts were mostly unsuccessful, as the generated images tended to vary from the original architectural image rather than providing a two-dimensional interpretation. Even when the participants successfully used the generated image, the results were only loosely related to the original images. The visual search was over at a certain point, and the participants employed modeling software (e.g., Revit) to produce a 3D model from one or more images.

Initially, the participants compared the work-in-progress 3D model to the AI images. However, this became less frequent as the modeling process progressed. Interestingly, most participants did not employ the AI models to explore the design during the modeling phase further.

Participants generally tended to remain faithful to the selected AI images, which we argue is a design fixation (see Figure 3). The clearer and more precise the images were, the more the participants replicated the design in the image rather than dealing with the complexity associated with transitioning it to 3D.

One participant was observed to provide a 3D section screenshot image to Dall-E and requested that the section be developed in a specific known architect's style. Following the development of a model, some participants employed AI models to enhance the renders of the 3D building model. This entailed modifications to the weather, expansion of the image, and the incorporation of local context. Additionally, one participant utilized the generated images as artistic components to decorate an interior space render. One participant chose to pass from using the AI and created their design independently instead. The participant explained that she was dissatisfied with all the outputs and, thus, rejected its involvement in the creative process.

DISCUSSION

This study of how architecture students use image-generating models for design revealed that the models are best used in the early ideation phase. Additionally, examining text prompts used in the process uncovered distinct patterns that contributed to creating architectural designs. Interestingly, even though students used the models with similar frequency, they repurposed the outputs from MidJourney significantly more (40% of the time compared to 23% for the rest). These findings help to further our understanding of the role of image-generating models in the architectural design process.

Our results have demonstrated the potential of generative models in computational creativity, specifically in architectural design. Utilizing combinatorial creativity, as suggested by Boden (2004), our results indicate that generative models can be used to search for ideas, precedents, and inspiration, as well as to merge concepts, such as two architectural styles, to produce creative outcomes (see Figure 4). This suggests that generative models can be used to explore various novel design solutions, thereby enabling architects to leverage computational creativity to explore new aesthetic possibilities.

Some results suggest that the use of AI models to create design images can be problematic, as it can lead to design fixation. Specifically, when participants used MidJourney, which output illustrations were more "artistic," the participants tended to remain more consistent with the generated images, copying the features precisely and producing a result very close to the picture (see Figure 3). In contrast, participants that used more models that made more realistic pictures, such as Dall-E and Stable Diffusion, demonstrated more creativity in creating architecture, as the design in the generated images was less clear. This suggests that participants may have become "fixated" on MidJourney's beautiful illustrations, being less critical, and did not make the necessary adjustments. These findings have important implications for using AI models in the design process, as they suggest that design fixation can occur.

Figure 3
Various examples of image to 3D conversion. On the left, MidJourney generated images, on the right, the produced BIM models using Revit.

Furthermore, the utilization of the models in the advanced phases of the design process has been limited in their efficacy. This was evident when participants attempted to use the models to develop designs. The participants tried to provide an image of the designs they wanted to develop with text prompts to improve the design or transform between architectural blueprints (e.g., plans, sections, or elevations), yet the results were frustrating. This is likely due to the random nature of the diffusion models, which can generate various images, but at the expense of continuity. These results hinder the application of these models for design iterations when the goal is to advance the design.

Figure 4
Example of combinational creativity using a prompt that merges Gaudi's with Zaha Hadid's style. Left images by Dall-E, right images by MidJourney.

In the context of generating blueprints, the models failed to adhere to engineering logic. The blueprints were unclear, spaces were not defined, openings and doors were missing, and sometimes there was a mix between plan and section views. Furthermore, when multiple plans, sections, and facades were included in a single image, no discernible engineering logic connected them. These models did not have the knowledge of architectural principles, which was partly demonstrated by previous works (Chaillou, 2020).

Nonetheless, the simplicity of using these models for architectural expression has the potential to revolutionize how clients communicate their ideas. By allowing non-professionals to create new and innovative images and designs, AI models can provide a platform for the visual articulation of ideas that have been difficult to express without graphical skills. However, to effectively utilize this tool, acquiring knowledge of architectural terminology, as identified in this study, and gaining experience in using the instrument is necessary. Additionally, architects can use AI models to generate multiple design alternatives to present to clients, thereby facilitating the exploration of design directions.

Limitations
The limitations of our study stem from its exploratory nature, focusing on architecture students in a studio setting, which may not accurately represent the experiences of professional architects, potentially influencing creative fixation.

CONCLUSION
In conclusion, our study explored the potential of image-generating AI models in architectural design by examining their use in a studio context. We analyzed model usage and architectural prompts identifying their structural anatomy, benefits, and limitations. Our findings show that AI models can inspire creativity and innovation. However, they face challenges like causing design fixation, having limited design development capability, and generating non-coherent blueprints. These challenges highlight the need for further research to improve AI models' effectiveness in architectural design.

Future research should focus on reducing creative fixation, enhancing AI models' applicability to advance a design by better understanding human intent, and improving architectural logic by developing AI models that generate coherent plans and corresponding sections. With these research directions, we aim to contribute to the ongoing discussion on AI's role in architecture and facilitate more effective integration of image-generating models in the architectural design process.

REFERENCES

Boden, M.A. (2003). The Creative Mind: Myths and Mechanisms (2nd ed.). Routledge. https://doi.org/10.4324/9780203508527https://doi.org/10.4324/9780203508527

Boim, A., Dortheimer, J. and Sprecher, A. (2022) 'A Machine-Learning Approach to Urban Design

Interventions In Non-Planned Settlements', in. CAADRIA 2022: Post-Carbon, Sydney, Australia, pp. 223–232. Available at: https://doi.org/10.52842/conf.caadria.2022.1.223.

Bolojan, D., Vermisso, E. and Yousif, S. (2022) 'Is Language All We Need? A Query Into Architectural Semantics Using a Multimodal Generative Workflow', in. CAADRIA 2022: Post-Carbon, Sydney, Australia, pp. 353–362. Available at: https://doi.org/10.52842/conf.caadria.2022.1.353.

Chaillou, S. (2020). ArchiGAN: Artificial Intelligence x Architecture. In: Yuan, P.F., Xie, M., Leach, N., Yao, J., Wang, X. (eds) Architectural Intelligence. Springer, Singapore. https://doi.org/10.1007/978-981-15-6568-7_8

Dortheimer, J. (2022) 'Collective Intelligence in Design Crowdsourcing', Mathematics, 10(4), p. 539. Available at: https://doi.org/10.3390/math10040539.https://doi.org/10.3390/math10040539.

Goodfellow, I.J. et al. (2014) 'Generative Adversarial Networks'. arXiv. Available at: http://arxiv.org/abs/1406.2661 (Accessed: 19 March 2023). (Accessed: 19 March 2023).

Epstein, Z., Schroeder, H. and Newman, D. (2022) 'When happy accidents spark creativity: Bringing collaborative speculation to life with generative AI'. arXiv. Available at: http://arxiv.org/abs/2206.00533 (Accessed: 20 January 2023). (Accessed: 20 January 2023).

Eroğlu, R. and Gül, L.F. (2022) 'Architectural Form Explorations through Generative Adversarial Networks - Predicting the potentials of StyleGAN', in. eCAADe 2022: Co-creating the Future - Inclusion in and through Design, Ghent, Belgium, pp. 575–582. Available at: https://doi.org/10.52842/conf.ecaade.2022.2.575.

Gmeiner, F., Yang, H., Yao, L., Holstein, K. and Martelaro, N., 2023. Exploring Challenges and Opportunities to Support Designers in Learning to Co-create with AI-based Manufacturing Design Tools. arXiv preprint arXiv:2303.00192.

Huang, J. et al. (2021) 'On GANs, NLP and Architecture: Combining Human and Machine Intelligences for the Generation and Evaluation of Meaningful Designs', Technology|Architecture + Design, 5(2), pp. 207–224. Available at: https://doi.org/10.1080/24751448.2021.1967060.

Liu, V. et al. (2022) '3DALL-E: Integrating Text-to-Image AI in 3D Design Workflows'. arXiv. Available at: http://arxiv.org/abs/2210.11603 (Accessed: 28 February 2023). (Accessed: 28 February 2023).

Maher, M.L., Poon, J. and Boulanger, S. (1996) 'Formalising Design Exploration as Co-Evolution', in J.S. Gero and F. Sudweeks (eds) Advances in Formal Design Methods for CAD. Boston, MA: Springer US, pp. 3–30. Available at: https://doi.org/10.1007/978-0-387-34925-1_1. .

Mansimov, E. et al. (2016) 'Generating Images from Captions with Attention'. arXiv. Available at: http://arxiv.org/abs/1511.02793 (Accessed: 20 January 2023). (Accessed: 20 January 2023).

Mrosla, L. and Koch, V. (2019) 'Quo vadis AI in Architecture?' Available at: https://doi.org/10.5151/proceedings-ecaadesigradi2019_302. .

Negroponte, N. (1972) The Architecture Machine: Toward a more human environment. 2. print. Cambridge, Mass.: MIT Pr.

Olteţeanu, A.-M. (2020) Cognition and the Creative Machine: Cognitive AI for Creative Problem Solving. Cham: Springer International Publishing. Available at: https://doi.org/10.1007/978-3-030-30322-8

Oppenlaender, J. (2022) 'A Taxonomy of Prompt Modifiers for Text-To-Image Generation'. arXiv.

Available at: http://arxiv.org/abs/2204.13988 (Accessed: 20 January 2023). (Accessed: 20 January 2023).
Picon, A. (2020) 'What About Humans? Artificial Intelligence in Architecture', in P.F. Yuan et al. (eds) Architectural Intelligence. Singapore: Springer Nature Singapore, pp. 15–29. Available at: https://doi.org/10.1007/978-981-15-6568-7_2.
Ramesh, A. et al. (2021) 'Zero-Shot Text-to-Image Generation'. arXiv. Available at: http://arxiv.org/abs/2102.12092 (Accessed: 20 January 2023).
Rombach, R. et al. (2022) 'High-Resolution Image Synthesis with Latent Diffusion Models', in 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA: IEEE, pp. 10674–10685. Available at: https://doi.org/10.1109/CVPR52688.2022.01042
Saharia, C. et al. (2022) 'Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding'. arXiv. Available at: http://arxiv.org/abs/2205.11487 (Accessed: 14 February 2023). (Accessed: 14 February 2023).
Silvestre, J., Ikeda, Y. and Guéna, F., 2016. Artificial imagination of architecture with deep convolutional neural network. Living Systems and Micro-utopias: Towards Continuous Designing, Melbourne; The Association for Computer-Aided Architectural Design Research in Asia (CAADRIA): Hong Kong, China, pp.881-890.
Simon, H.A., 1969. The sciences of the artificial Cambridge, MA: MIT Press
Yousif, S., Vermisso, E. (2022), Towards AI-Assisted Design Workflow for an expanded design space.
Zhang, G. et al. (2021) 'A cautionary tale about the impact of AI on human design teams', Design Studies, 72, p. 100990. Available at: https://doi.org/10.1016/j.destud.2021.100990
Zheng, H. et al. (2020) 'Apartment Floor Plans Generation via Generative Adversarial Networks', in. CAADRIA 2020: RE:Anthropocene, Bangkok, Thailand, pp. 599–608. Available at: https://doi.org/10.52842/conf.caadria.2020.2.599..
Zhu, J.-Y. et al. (2018) 'Toward Multimodal Image-to-Image Translation'. arXiv. Available at: http://arxiv.org/abs/1711.11586 (Accessed: 1 February 2023). (Accessed: 1 February 2023).