Spatial Learning for Functional Hybridisation

Critical methodology for generative architectural suggestions in concept design phase

Roberto Pasquale Romeo[1], Federica Pradella[2], Giorgio Castellano[3], Fabio Bazzucchi[4], Ingrid Maria Paoletti[5]
[1,2,3,5]Politecnico di Milano [4]Politecnico di Torino
[1]pasqualeroberto.romeo@mail.polimi.it,
[2,3,5]{federica.pradella|giorgio.castellano|ingrid.paoletti}@polimi.it
[4]fabio.bazzucchi@polito.it

The research aims to define a methodology to support the creative process in the early architectural design phase through Artificial Intelligence (AI), focusing on spatial learning and generative typology hybridization. It specifically investigates how AI can interpret and suggests internal spatial configurations of buildings by leveraging Generative Adversarial Networks (GANs). In recent years, the rise of AI-generated art has challenged traditional notions of creativity, and its impact is now increasingly visible in architectural design processes. The proposed approach integrates AI into design by exploring the latent space of StyleGAN, training AI on building sections' data and reconstructing 3D forms through point clouds. This offers a hybrid workflow between AI and computational modelling to generate three-dimensional architectural forms. As a design tool for architects, this workflow enhances spatial understanding of a building, analysing and exploring spaces of different typologies, looking for potential spatial hybridisation in design. The methodology is grounded in a critical reassessment of AI's role in design creativity, not as a replacement for human intuition but as a collaborator capable of producing unexpected spatial interpolations. The study highlights how hallucination, often considered a flaw in generative AI, can become a productive tool for design speculation, enabling new forms of typological hybridisation. This research contributes to redefining the epistemology of architectural form-making in the age of intelligent systems, opening a path toward collaborative creativity between human designers and artificial agents.

Keywords: *Creativity Augmentation, Artificial Intelligence, Architectural Design methodology, Spatial Understanding, GAN*

INTRODUCTION

The evolving relationship between Architecture and Artificial Intelligence (AI) is at the core of this research. Technological innovation has consistently shaped the architectural discipline, and the growing diffusion of AI-based tools, particularly generative models, now represents a potential paradigm shift, challenging conventional design thinking and production. This research investigates the potential of a new approach to early-stage design by developing a methodology with AI at its core, learning and proposing interior spatial configurations. The study explores how AI can become a generative partner in architectural thinking by leveraging the latent space of StyleGANs and reconstructing 3D spatial suggestions from 2D drawing data.

Research motivation

In recent years, architectural practices have increasingly engaged with AI applications, mainly through generative AI platforms (deep learning text-to-image model, i.e. (Midjourney, no date), (Stability AI, no date)), for the generation of early-stage concept phase ideas and visualizations.

Although diffusion models (Koehler, 2023) have become increasingly pervasive in the concept phase of the design workflow, new approaches are emerging. These include diffusion pipeline with ControlNet (Chen *et al.*, 2024), Graph Neural Networks (Zhao *et al.*, 2023) and Transformative-based Generative Networks (Hudson and Zitnick, 2021), each exploring alternatives methods of generating and manipulating architectural data.

StyleGAN is one of the first models applied to architectural concept generation and artistic experimentation (Berryman, 2024). Thanks to its fast execution, latent space structured exploration, and formal consistency, it enables smooth interpolation between design variants, making it a promising tool for morphological analysis.

As investigated by Matias del Campo (del Campo, 2021), it can be effectively trained on style-specific architectural datasets, ensuring coherent and recognizable outputs.

The architectural potential of this technology was first extensively explored by Stanislas Chaillou (Chaillou, 2020), who demonstrated how neural networks could translate building footprints into interior spatial layouts through a sequence of image-to-image transformation. However, moving from two-dimensional representation to full three-dimensional architectural models remains a complex issue, mainly due to the high complexity of implementation and the computational power required to process 3D information (Leach, 2022). Whether some attempts succeeded in generating three-dimensional objects (Stability AI, no date), such results could only give the envelope without any topological interior information, limiting their practical application in architectural design.

Research objectives

This research aims to support the creative process in the early architectural design phase through AI, leveraging StyleGAN, to generate spatially coherent and topologically expressive 3D outputs. The proposed methodology combines AI and computational tools into a 2D-to-3D generative workflow. In this system, the training data consists of slices, in this case sections, from a large dataset of architectural objects (Zhang, 2019a), enabling AI to learn. Generating a three-dimensional object relies on the assembly sequence of sections' images converted into point clouds (Kuptsova, 2022), (Bank et al., 2022). This methodology is intended to assist designers during the conceptual design phase by enhancing spatial and formal exploration through AI-generated suggestions.

THEORETICAL BACKGROUND

Early intersections between architecture and AI can be traced to the work of (Campo et al., 2019) and (Chaillou, 2020), who explored the design opportunities enabled by the Generative Adversarial Networks (GANs) model. Similarly, (Zhang, 2019) introduced a slicing-based methodology for developing 3D geometry. GANs were employed not to generate isolated images but to produce a sequence of interrelated slices navigated through latent space. The latent space (Nature Machine Intelligence, 2020), often described as the black box of neural networks, becomes central to this research. Rather than generating a static, singular image, the process generates multiple sequential images through a vectorial interpolation. This conceptual journey in the brain of the model is allowed by a specific model of GAN, known as the StyleGAN.

GANs and StyleGANs

Generative Adversarial Networks (GANs), first introduced by Ian Goodfellow in 2014 (Miller, 2019), have undergone significant advancements in recent years, setting the standard of image generation regarding precision and resolution.

As Neil Leach notes, "*There is no way to tell how or why they are performing certain operations. In this case, they are a little different to intuition. Both GANs and human intuition are black boxes*" (Leach, 2022). This unpredictable nature makes GANs particularly suitable for early-stage experimental design, where indeterminacy can be assessed (Prix et al., 2022).

Although the resolution of images produced by GANs have improved significantly over time (Brock, Donahue and Simonyan, 2018), the generator has continued to operate as black box, with a lack of understanding of various aspects of the image synthesis process. This is relevant, especially within architectural workflows, where GAN-generated images' limited controllability and spatial depth pose challenges (Gupta et al., 2024). In addition, the properties of the latent space are poorly understood, and the commonly demonstrated latent space interpolations (Choi *et al.*, 2017) provide no quantitative way to compare different generators against each other.

In response to these limitations, a breakthrough came with the development of StyleGAN (Karras, Laine and Aila, 2018), which introduced a redesigned generator architecture that enables more nuanced control over the synthesis process. Instead of relying on a traditional latent vector input, StyleGAN begins from a learned constant and modulates the output through style vectors applied at each convolutional layer. This allows for hierarchical control of image features across scales ranging from global structures to fine-grained textures. A significant advantage of StyleGAN is that it does not require labelled datasets, which are typically time-consuming to prepare. The model autonomously extracts relevant features from the training data. By training on curated datasets, StyleGAN can produce images that reflect the formal language and spatial logic embedded in the input references (Karras *et al.*, 2020).

Hallucination and AI

Creativity has been intended as a unique and special human ability rooted in emotion, intuition and unconscious associations (Gu *et al.*, 2018). Recent developments in art have shown that AI can generate art indistinguishable from that made by humans (Obvious Collective, 2018) (Klingemann, 2017). Within human and machine creativity boundaries, generating unexpected or misleading information based on internal statistical inferences is known as “hallucination” (del Campo and Leach, 2022). This phenomenon has been investigated by Refik Anadol, who in 2018 was commissioned to produce a data sculpture for the LA Philharmonic. Based on the building’s historical archive, the project gave a memorial heritage to an AI, questioning whether a machine could also elaborate memories (Leach, 2022).

In Artificial Neural Networks (ANNs), this elaborative capacity is within the latent space, a virtual multi-dimensional space which contains a representation of the data elaborated during the training. Each dimension encodes a feature not predefined by the user but shaped through the model’s learning process (Chaillou, 2022), meaning that the ANN recognizes features in the data that the user did not foresee. This is the “latent” part of the process. In this space, proximity between data points reflects similarity: elements that look alike are mapped closer together. Moreover, latent spaces are often topologically consistent, meaning that any two points can be interpolated to create new data samples that blend characteristics of both

(Radford, Metz and Chintala, 2015). Each image corresponds to a coordinate in the latent space, defined by the number of embedded features or its dimensionality.

METHODOLOGY

This methodology is designed to integrate AI-generated creativity within the architectural design process. The primary input consists of architectural section drawings extracted from a custom 3D dataset. Following a training phase, new images are generated, subsequently assembled and translated into three-dimensional representations. This workflow integrates AI, Rhinoceros, and Grasshopper to create the dataset and to translate and assemble the generated images in a three-dimensional visualization. The aim is to develop a generative model that displays typological hybridization, inviting reflections on the interior space. Sections' images are preferred over plans since they provide a deeper spatial understanding for reconstructing 3D forms from 2D data and capturing programmatic distribution across spaces. The three-dimensional reconstruction is implemented by point cloud (Bank *et al.*, 2022).

The sections are extracted from two different architectural typologies: public residential buildings from contemporary Italian architecture (from 2000 onward) and temporary pavilions (mainly from International Expositions held in 2010, 2015, and 2020). This selection ensured a homogenous dataset, avoiding different countries' building codes and regulatory frameworks.

Dataset acquisition and preparation

To generate the dataset (Figure 1), 28 residential buildings have been selected among towers, multi-apartments, courtyards, and 28 pavilions. Buildings have been modelled using Rhinoceros with 11 colours (following the HTML hexadecimal code), which represent different functions and elements:

- black (#ff0000) for walls and slabs
- purple (#ff00ff) for balconies and loggias
- red (#ff0000) for dining rooms
- violet (#9144d9) for kitchens
- blue (#0000ff) for bedrooms
- cyan (#00ffff) for bathrooms and toilets
- dark green (#008f00) for corridors
- yellow (#ffff00) for secondary spaces such as storage spaces, studios, meeting rooms
- light blue (#cdceff) for vertical distributions of stairs and lifts
- light green (#80ff00) for greenery and vegetation
- pink (#ffc2c2) for common spaces.

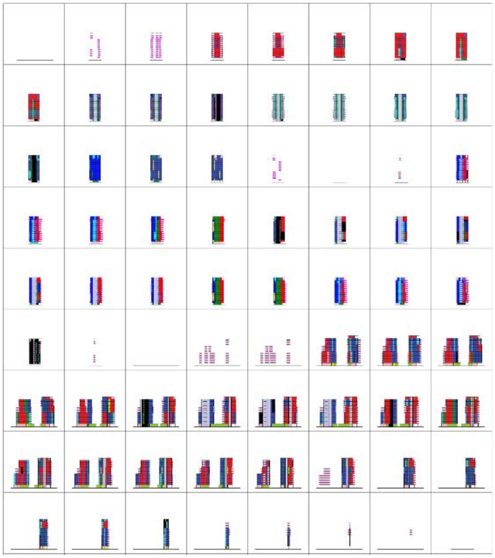

Figure 1
Dataset sample

The sections have been extracted as jpeg images from the modelled buildings through a slicing script in Grasshopper, with a step of 1 meter. Redundant data have been removed, and a final dataset consisting of 3952 images was created.

Secondary spaces, vertical distributions, common spaces, and greenery are common between the two typologies of buildings data, allowing sections to be found next to each other inside the latent space. It is important to highlight that the process does not directly implement all the variables which may define the quality or the dimensions of space; such features are supposed to be implicitly learnt from the model during the training.

Technical setup and StyleGAN training

The StyleGAN version employed is StyleGAN2-ADA-Pytorch (Karras *et al.*, 2020) which was selected for its ability to generate effective results even when working with limited datasets. The training was conducted on a workstation with an NVIDIA RTX 6000 Ada GPU.

StyleGANs need a wide range of data to learn from and to enrich the latent space composition. During training, the model processes the overall dataset through multiple computing cycles, called epochs of training. After each training epoch, the model undergoes backpropagation, a self-corrective algorithm used in deep neural networks to adjust internal weights based on error gradients (Rumelhart, Hinton and Williams, 1986). Achieving optimal results typically requires multiple epochs, with each epoch's duration depending on the input data's complexity and characteristics. For image-based datasets, resolution is critical in training time and output fidelity. This study identified a 512 × 512 pixels resolution as a practical compromise between visual quality and computational efficiency.

A pre-trained model has been used to accelerate training and improve output consistency, leveraging transfer learning (Mahajan *et al.*, 2018). Specifically, the model was initialized with the FFHQ pre-trained model, known for the good-quality generative results, proposed by Nvidia (Karras *et al.*, 2020), (NVIDIA, 2021). To preserve control over the training process and avoid unintended variation, no data augmentation techniques, such as cropping, rotation, flipping, or brightness alterations, were applied (Shorten and Khoshgoftaar, 2019).

Spatial generation process

The generative process in StyleGAN is performed through the inference, which consists of specifying a numeric value, called a seed, corresponding to an array of coordinates that localise data (in this case, an image) within the latent space (Goodfellow *et al.*, 2014).

Multiple seeds can be specified, enabling interpolation and generating a continuous series of images that transition between two selected points in the latent space. For the seeds' selection, two criteria have been explored:

1. Manual selection. Due to the high dimensionality of latent space, it is impossible to locate all sections, a sample of 500 images has been generated, analysed, and categorised according to their topology. Representative seeds are then selected to explore variations across building types, such as multi-apartment, towers, and courtyard buildings.
2. Random selection. In this streamlined approach, seed numbers have been randomly generated and used for interpolation without any preliminary selection of the sections, accelerating the iteration and allowing the analysis of a broader range of configurations.

While manual selection offers interpretability, this research considers the random method, which avoids manual filtering and enhances the model's capacity to generate unusual and unexpected architectural suggestions.

The outputs were evaluated using short interpolation videos, visualising the transformation between generated sections and quickly revealing the embedded spatial characteristics. Once promising outputs were identified, individual frames from the video sequences were assembled into a point cloud for

3D reconstruction. These videos serve as an effective preliminary filter, offering an early visual approximation of the point cloud and helping to determine whether the generated result is architecturally relevant.

Reconstruction process

The conversion and assembly of the generated images were carried out using a custom Grasshopper algorithm. The array of images composes a linear point cloud representation measuring approximately 50 meters in length, a dimension considered appropriate for the scale of architectural design proposals. In addition, part of the Grasshopper script included densification of the point cloud, aiming at enriching the spatial detail and quality of the output.

ANALYSIS & RESULTS

The following section presents the analysis of the dataset and the outcomes of the AI-driven generative process, highlighting spatial coherence and design potential.

Dataset evaluation

The training dataset contains all the features the model is supposed to learn, and it was carefully curated to ensure meaning content, avoiding redundancy. Excluding the building structure, the dataset is slightly unbalanced toward the common spaces, followed by dining rooms and bedrooms. This asymmetry was intentionally preserved to guide the model toward generating outputs with a stronger presence of shared spatial functions.

In addition, the image variety has been assessed and visualized through Principal Component Analysis (PCA) algorithm, a dimensionality reduction technique that project high-dimensional data (like images described by many pixel values or features) and projects it into a lower-dimensional space while retaining the most relevant features (Permanajati, 2022). The complexity of latent space is measured through the entropy parameter, as seen in (Papia, Kondi and Constantoudis, 2023) (Chen *et al.*, 2019). As shown in Figure 2, the dataset exhibits a well-distributed range of visual complexity. A dense central cluster corresponds to simple, monochromatic images with low entropy (blue values), while increasingly dispersed data points indicate more intricate and colourful configurations with higher entropy (yellow values). Finally, the training validation (Figure 3) was performed on a collection of 4000 generated images to understand to what extent the model has learned the original features of the dataset. The distribution of the generated outputs appears more concentrated and does not reach the same level of complexity observed in specific training images. However, this reduced complexity is balanced by a broader distribution of variation across the generated samples, resulting in more diverse and often unconventional formal outcomes.

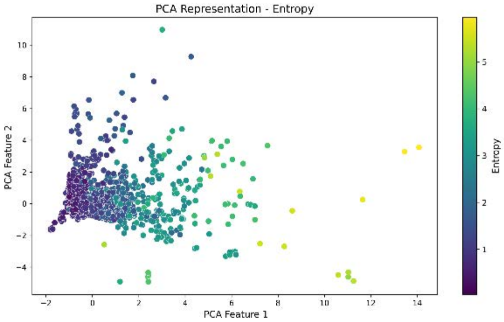

Figure 2
PCA dataset

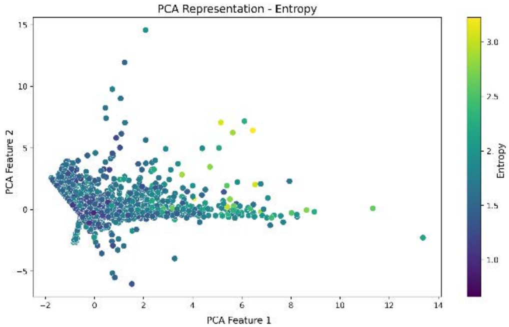

Figure 3
PCA trained data

Results

A series of generations was produced using the manual selection criteria (Figure 4) and the random selection criteria (Figure 5). In both cases, space organization was qualitatively respected.

From the point clouds, it is observable that balconies and loggias consistently appear on the external edges of the building. Floor plans analysis reveals that ground floors tend to reproduce typical features of residential typologies, with common spaces, secondary rooms, corridors, toilets, vertical distribution, and vegetation. The functional distribution of upper floors also aligning with usual building residential layouts.

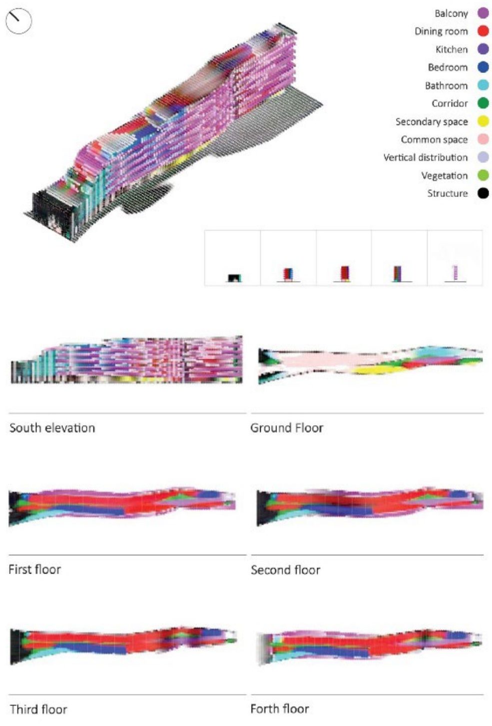

Figure 4
Final output
Manual selection

The outputs generated following the randomized approach (Figure 5) often present a hybridization of spaces coming from both public and private spaces. The frequent recurrence of common spaces in these outputs suggests that shared programmatic elements were among the most strongly learned features by the model.

Overall, the 2D-to-3D process employed is efficient and competitive in respect to its capacity to produce architecturally coherent results, providing internal spatial logic and shape of the generated object.

Operational workflow

The practical usage of this methodology can be divided into three main steps: image generation, selection of the output, and 3D reconstruction. Once the dataset has been modelled and thoroughly trained, the generation process is performed through Python scripts, which load the model and perform the inference supporting relatively fast iterations. After reaching the desired number of generations, a prior selection of the outputs is recommended since the reconstruction step may take minutes according to the hardware employed. This selection can be effectively evaluated by reviewing the interpolation video, which offers a quick yet insightful overview of the latent features presents in the generated samples. Finally, the selected images are imported into Grasshopper, assembled and converted into a point cloud format. This step enables spatial reconstruction and further integration into 3D design workflows.

OPEN ISSUES AND FUTURE DIRECTIONS

Some critical aspects currently limit the applicability of the proposed methodology. Firstly, the dataset production relies on manually constructed 3D models with a prior architectural projects' information retrieval, this step is time-consuming and may restrict scalability. Secondly, the format of the output presents a challenge.

Point clouds proved to be the most suitable solution for research purposes; however, they do not represent a conventional 3D format for architectural design, which is mostly based on NuRBS or MESH modelling. Finally, the generations tend to produce elongated or distorted spatial forms. A formal normalization step could help regulate these proportions and improve the usability of interior layouts for architectural consideration.

Potential improvements and next steps

Multiple technical improvements can be made to the workflow to obtain and construct the data. Currently, model training and generation are performed outside the 3D environment. Indeed, implementing the generative step as part of the Grasshopper script can further improve in workflow efficiency and timing.

In addition, given the need for compatible 3D formats, implementing a pipeline that generates shapes from clustered points (in this case, clustered by colours) may offer a designer-friendly format while facilitating compatibility with conventional modelling tools.

Furthermore, improvements can concern the training dataset. Defining a custom dataset inevitably involves a set of subjective choices that embed the author's perspective into the learning process. Optimizing the labelling procedure by introducing more systematic or objective classifications could reduce this bias and increase the dataset's neutrality. For example, including labels related to structural or material properties might enrich the generative potential while anchoring the output in more consistent architectural data. This refinement would make the process more robust and less susceptible to individual interpretive influence.

Regarding the generation process, a potential improvement involves managing the clustering structure within the latent space, allowing the latent exploration to be coherent with design needs. Indeed, after the training process, images are organized inside the latent space in clusters based on feature similarities.

Accessing and manipulating this clustering system would enable latent space exploration to be guided by specific spatial features, moving beyond purely formal outcomes and reducing reliance on random processes.

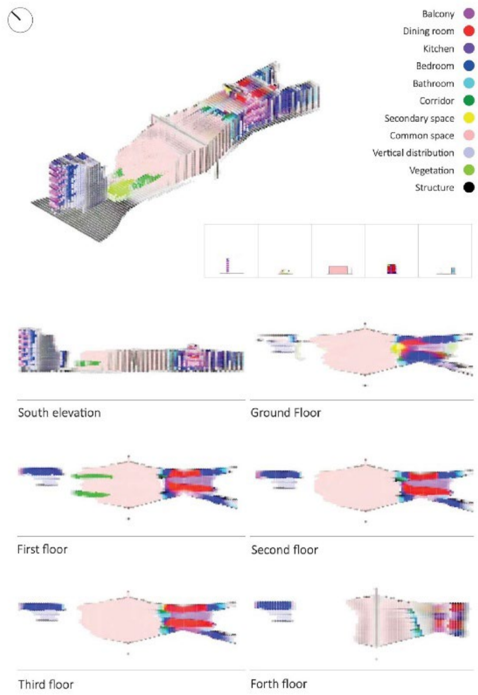

Figure 5
Final output
Random selection

Various parameters can be adjusted to influence the output of the generation process. One of the most significant is the number of seeds used for interpolation. A higher number of seeds tends to yield higher spatial complexity. On the other hand, using fewer seeds may lead to more linear and predictable outcomes, although it can also amplify hallucinations. These anomalies are particularly evident in low-density interpolations, where the model fills in gaps with configurations from underrepresented regions of the latent space.

CONCLUSIONS

Overall, this research highlights AI’s capacity to translate a latent space “hallucination” into tangible spatial configurations, fostering architectural design thinking by introducing new possibilities into the design process.

The 2D-to-3D approach preserves functional intelligibility, mapping out spaces such as balconies, corridors, and communal areas in ways that resonate with established architectural topology and expand its traditional boundaries. At the same time, the use of randomized seed interpolation highlights AI’s expansive generative potential, enabling designers to investigate unexpected hybridizations and novel typological configurations.

By continuing to refine how models learn and represent internal spatiality, architects and researchers can envision an era of collaborative creativity, where AI-driven experimentation complements human intuition in shaping new and unexpected built environments. It invites the discipline to reconceive space, function, and form boundaries.

REFERENCES

Bank, M., Sandor, V., Schinegger, K. and Rutzinger, S. (2022) ‘Learning Spatiality - A GAN method for designing architectural models through labelled sections’ in. *Education and Research in Computer Aided Architectural Design in Europe (eCAADe)*, pp. 611–619. https://doi.org/10.52842/conf.ecaade.2022.2.611.

Berryman, J. (2024) ‘Creativity and Style in GAN and AI Art: Some Art-historical Reflections’ *Philosophy and Technology*, 37(2). https://doi.org/10.1007/s13347-024-00746-8.

Brock, A., Donahue, J. and Simonyan, K. (2018) 'Large scale GAN training for high fidelity natural image synthesis' *arXiv (Cornell University)* [Preprint]. https://doi.org/10.48550/arxiv.1809.11096.

del Campo, M. (2021) ‘Architecture, Language, and AI. Language, Attentional Generative Adversarial Networks (AttnGAN) and Architecture Design’ in, pp. 211–220. https://doi.org/10.52842/conf.caadria.2021.1.211.

del Campo, M. and Leach, N. (2022) ‘Can Machines Hallucinate Architecture? AI as Design Method’ *Architectural Design*, 92(3), pp. 6–13. https://doi.org/10.1002/ad.2807.

Chaillou, S. (2020) ‘ArchiGAN: Artificial Intelligence x Architecture’ in P.F. Yuan et al. (eds) *Architectural Intelligence*. Springer. Singapore: Springer.

Chaillou, S. (2022) *Artificial Intelligence and Architecture. From Research to Practice*. Birkhäuser.

Chen, J., Zheng, X., Shao, Z., Ruan, M., Li, H., Zheng, D., and Liang, Y. (2024) ‘Creative interior design matching the indoor structure generated through diffusion model with an improved control network’ *Frontiers of Architectural Research* [Preprint]. https://doi.org/10.1016/j.foar.2024.08.003.

Chen, X., Zhang, Q., Lin, M., Yang, G., and He, C. (2019) ‘No-reference color image quality assessment: from entropy to perceptual quality’, *EURASIP Journal on Image and Video Processing*, (1), p. 77.

Choi, Y., Choi, M., Kim, M., Ha, J.-W., Kim, S. and Choo, J. (2017) 'STARGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation,' *arXiv (Cornell University)* [Preprint]. https://doi.org/10.48550/arxiv.1711.09020.

Goodfellow, I.J. *et al.* (2014) 'Generative adversarial networks' *arXiv (Cornell University)* [Preprint]. https://doi.org/10.48550/arxiv.1406.2661.

Gu, S., Gao, M., Yan, Y., Wang, F., Tang, Y.-y., and Huang, J.H. (2018) 'The Neural Mechanism Underlying Cognitive and Emotional Processes in Creativity' *Frontiers in Psychology*. Frontiers Media S.A.

Hudson, D.A. and Zitnick, C.L. (2021) 'Generative adversarial transformers' *arXiv (Cornell University)* [Preprint]. https://doi.org/10.48550/arxiv.2103.01209.

Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., and Aila, T. (2020) 'Training Generative Adversarial Networks with Limited Data' *arXiv (Cornell University)* [Preprint]. https://doi.org/10.48550/arxiv.2006.06676.

Karras, T., Laine, S. and Aila, T. (2018) 'A Style-Based generator architecture for generative adversarial networks' *arXiv (Cornell University)* [Preprint]. https://doi.org/10.48550/arxiv.1812.04948.

Klingemann, M. (2017) *The Butcher's Son*. Available at: https://www.electricartefacts.art/artwork/mario-klingemann-the-butchers-son-special-edition

Koehler, D. (2023) 'More than anything: Advocating for synthetic architectures within large-scale language-image models' *International Journal of Architectural Computing*, 21(2), pp. 242–255.

Leach, N. (2022) *Architecture in the Age of Artificial Intelligence, An Introduction to AI for Architects*. Bloomsbury.

Mahajan, D. *et al.* (2018) 'Exploring the limits of weakly supervised pretraining' *arXiv (Cornell University)* [Preprint]. https://doi.org/10.48550/arxiv.1805.00932.

Midjourney, Inc. (n.d.) *Midjourney*. Available at: https://www.midjourney.com/home

Miller, A. (2019) *The Artist in the Machine: The World of AI-Powered Creativity*. MIT Press.

Nature Machine Intelligence (2020) 'Into the latent space' *Nature Machine Intelligence*, 2(3), pp. 151–151.

NVIDIA (2021) *Flickr-Faces-HQ Dataset (FFHQ)*. Available at: https://github.com/NVlabs/ffhq-dataset

Obvious Collective (2018) *Edmond de Belamy*. Available at: https://en.wikipedia.org/wiki/Edmond_de_Belamy

Papia, E.-M., Kondi, A., and Constantoudis, V. (2023) 'Entropy and complexity analysis of AI-generated and human-made paintings' *Chaos, Solitons & Fractals*, 170, p. 113385. https://doi.org/10.1016/j.chaos.2023.113385.

Permanajati, A.R. (2022) *Principal Component Analysis (PCA)*. Available at: https://docs.opencv.org/4.x/d1/dee/tutorial_introduction_to_pca.html

Radford, A., Metz, L. and Chintala, S. (2015) 'Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks' *arXiv (Cornell University)* [Preprint]. https://doi.org/10.48550/arxiv.1511.06434.

Rumelhart, D.E., Hinton, G.E. and Williams, R.J. (1986) 'Learning representations by back-propagating errors' *Nature*, 323(6088), pp. 533–536. https://doi.org/10.1038/323533a0.

Shorten, C. and Khoshgoftaar, T.M. (2019) 'A survey on Image Data Augmentation for Deep Learning' *Journal of Big Data*, 6(1). https://doi.org/10.1186/s40537-019-0197-0.

Stability AI (n.d.) *Stable Diffusion*. Available at: https://stability.ai/news/stable-zero123-3d-generation

Zhang, H. (2019) '3D Model Generation on Architectural Plan and Section Training through Machine Learning' *Technologies*, 7(4), p. 82. https://doi.org/10.3390/technologies7040082.

Zhao, P., Liao, W., Huang, Y., and Lu, X. (2023) 'Intelligent design of shear wall layout based on graph neural networks' *Advanced Engineering Informatics*, 55. https://doi.org/10.1016/j.aei.2023.101886.