

The Art of Storytelling Redefined: Artificial Intelligence and the Future of Architectural Visualizations

Styliani (Stella) Salta, PhDc^{1a}, Miltiadis Katsaros, Prof. As.¹

¹ School of Architecture, National Technical University of Athens

Keywords: Artificial Intelligence, Creativity, Architectural Design, Text-to-image generators, Visualizations

archiDOCT

Vol. 12, Issue 1, 2024

This paper explores the transformative influence of Text-to-Image (TTI) generative artificial intelligence (AI) on architectural visualizations, posing a critical inquiry into whether it serves as an optimized visualization tool or signals a paradigm shift in architectural design methodologies. Tracing the evolutionary path of TTI models—from early generative adversarial networks to contemporary diffusion models—the study underscores their integration into architectural workflows. Emphasizing the synergy between human creativity and AI insights, the research advocates for a redefined approach to conceptualization, self-knowledge, and collective imagination in design. Through iterative explorations, architects engage in a dialogue with TTI systems, fostering self-discovery and refining design identity. The paper further explores TTI's potential in stimulating creativity, drawing parallels to surrealist collage practices, and challenging designers to venture into unexplored conceptual realms. The study concludes by highlighting the hyper-dimensional nature of Text-To-Image generation, while unraveling the intricate web of authorship in the AI creative process, recognizing the intertwined roles of humans, engineers, algorithms and vast datasets. While emphasizing the imperative need for ethical frameworks, the paper acknowledges that Text-to-Image technology has the potential to stimulate architectural imagination in unprecedented directions, functioning as a catalyst for creativity rather than merely serving as a visualization tool.

1. Introduction

Text-to-Image (TTI) generators, fueled by generative artificial intelligence, have become a focal point of fascination across various creative disciplines (Oppenlaender, 2022). These innovative tools possess the capability to translate textual descriptions directly into visual representations, ranging from realistic scenes to imaginative concepts. This technological advancement has sparked significant curiosity, particularly amongst the field of design and architecture, signaling a departure from conventional visualization techniques.

The allure of Text-to-Image generators lies not only in their ability to bring textual ideas to life but also in the profound questions they raise about the creative process. Beyond serving as mere tools for visualizing predefined concepts, these generators are positioned as potential catalysts, influencing not only the visual aspects of design but also the very methodology and cognitive processes employed by architects. This explorative study lays the groundwork for a deeper examination of how text-to-image generators may reshape workflows, narratives and ways of thinking in early stages of conceptual design.

^a Styliani (Stella) Salta is an architect and computational designer. Her research focuses on the confluence of Architecture and Artificial Intelligence, concentrating on generative design, and the strategic integration of data science and machine learning methodologies within the realms of Architecture and Digital Arts.

She is actively engaged in the fields of digital fabrication and data-driven design, challenging conventional anthropocentric paradigms of creation. Her work reflects a dedication to exploring speculative notions of design, embracing self-organization practices, and tapping into the creative potential of non-human entities. Through these approaches, she seeks to amplify collective creativity, expand imaginative capacities, and refine problem-solving methodologies.

She graduated with a diploma in Architectural Engineering from the Polytechnic School of Aristotle University of Thessaloniki (Greece) in 2018. Her postgraduate studies include an MArch from NTUA School of Architecture and an MSc in Data Science and Machine Learning from NTUA School of Electrical and Computer Engineering. Currently, she is a PhD candidate at the National Technical University of Athens, School of Architecture, actively contributing to the evolving discourse at the intersection of architecture and advanced technologies.

2. Objectives

Departing from the main research question, discerning whether text-to-image generators represent a new visualization tool or signal a paradigm shift in design methodologies, there are three main objectives to be investigated when looking into the potential of these tools:

- **Redefining Conceptualization**

Can Text-to-Image generators transcend their conventional role as visualization tools and actively contribute to the conceptualization process, reshaping architects' approaches to conceiving and formulating design ideas?

- **Unleashing Self Knowledge**

In what ways do Text-to-Image generators serve as catalysts for self-discovery in architectural design? How do they enable architects to explore their creative potential and deepen their understanding of the intricacies of their own design thinking through iterative engagements?

- **Acknowledging the Influence of Collective Imagination**

How does the collaborative nature of creative thoughts, facilitated by Text-to-Image models trained on extensive datasets of various creators' works, contribute to fostering creativity within distributed cognition? To what extent does this process form the foundation for innovative architectural expressions by tapping into a shared creative consciousness?

Looking into the aforementioned objectives, this study proposes an evaluation framework for practitioners and creatives, positioning text-to-image generators as tools for concept articulation. The main parts of this framework are:

- **Iterative Explorations**

Text-to-image generators streamline the iterative design process, empowering designers to experiment and refine their ideas effortlessly. The emphasis is on reducing effort while facilitating idea iteration and maintaining the quality of design evolution.

- **Workflow Transformation**

The study suggests mapping of the design process to understand how these tools reshape workflows in architectural practice. The pivotal shift occurs as they enable the exploration of concepts through high-definition visualizations, advancing photorealistic presentations in the workflow beyond their current stage and eliminating the need for preliminary design, time investment and effort.

- **Creativity and Concept Making**

Exploring the potential of these tools to foster genuine creativity. Drawing inspiration from surrealist collage techniques, to examine unconventional cases that challenge architects to think creatively. The aim is not merely rapid architectural rendering but leveraging the TTI generators to enhance creative thinking and conceive innovative spatialities.

3. Text-to-Image AI

The integration of such tools into design workflows necessitates a fundamental understanding of their inner logic and evolution in the field of artificial intelligence (AI). Generative AI is a subfield within the broader realm of artificial intelligence that focuses on the development of systems capable of creating, generating, or producing new content autonomously. Unlike traditional AI approaches that often rely on explicit programming and rule-based systems, generative AI leverages advanced machine learning techniques, to enable systems to learn and understand patterns within vast amounts of data – or in a more precise description, built semi-reliable statistical correlations of information relationships found in the given data (Bernstein, 2022). The hallmark of generative AI is its ability to generate novel outputs, such as images, text and multimedia by “learning” from those datasets. (Goodfellow et al., 2020)

A demonstration of generative AI's prowess is witnessed in text-to-image (TTI) generation. TTI refers to a technology that involves the generation of visual content, such as images or graphics, based on textual descriptions or prompts. TTI operates by utilizing advanced deep learning models, trained on extensive datasets containing pairs of text and corresponding images, enabling them to learn the intricate relationship between textual descriptions and visual elements (Chaillou, 2022, p. 30). As a result, Text-To-Image systems can automatically generate high-fidelity images that closely match complex text descriptions. The synergy between generative models and large language models has played a pivotal role in enhancing the performance of TTI generation (Koh et al., 2023), often blurring the lines between human-created and AI-generated visual content. The most effective TTI algorithms have generally been trained on massive amounts of image and text data scraped from the web. Some of the most popular TTI models available online include: DALL-E by OpenAI, Imagen by Google Brain, Stable Diffusion by StabilityAI and Midjourney.

4. Brief history of TTI evolution

The evolution of TTI models is marked by the advent of deep learning in image generation around 2015 (Zhang et al., 2023). Generative adversarial networks (GANs) served as the primary backbone for TTI models, efficiently producing images through a generator-discriminator structure. GANs operate in tandem. The generator creates new data samples mimicking the training data, while the discriminator distinguishes between real and generated data. These networks undergo joint training in a zero-sum game, where the generator aims to create increasingly realistic samples, and the discriminator identifies fabricated ones. This process continues until the generator achieves the ability to generate data samples indistinguishable from authentic ones (Goodfellow et al., 2020). Through the collaborative training of these networks, GANs can develop the capability to generate high-quality images closely resembling those found in the real world.

Over time, researchers explored various architectures for text-to-image generation, leading to the emergence of autoregressive models and later on the diffusion model. Autoregressive models, also belonging in the domain of deep learning, are a category of models useful for generating images from text. These models operate by predicting each pixel in the image individually, using information from the previous pixels. Essentially, the model sequentially generates the image, with each pixel being produced based on the information gleaned from its preceding pixels (Shih et al., 2022). This approach mirrors the way humans draw images, initiating the process with a basic sketch and progressively adding finer details.

The latest approach involves diffusion models. They operate through an iterative diffusion process applied to a noise vector, progressively shaping this noise into a coherent image. The dynamics of this diffusion process are governed by adjustable parameters fine-tuned during training to ensure the production of high-quality images (Saharia et al., 2022). This methodology allows for a controlled and gradual transformation, wherein the noise vector evolves into a visually meaningful representation, thereby bridging the gap between textual information and visual output.

The ongoing exploration of different TTI model types, considering factors like model size, time efficiency, and image fidelity, illustrates a dynamic and competitive landscape in the development of AI-generated images. This journey has made TTI generation a focal point of discussion within the broader field of AI generated content. Despite the diverse models and methodologies employed by the research community, the user interface designed for interaction remains exceptionally simple. Users engaging with TTI models are not required to possess knowledge or comprehension of the underlying workings of the aforementioned algorithmic models, nor do they need any drawing or artisanal skills. From the user's standpoint, all that is essential is a prompt—a concise textual description of the envisioned image outcome. The ease of use has contributed to the widespread popularity of TTI across diverse users and disciplines, prompting exploration and integration attempts in various fields, as these systems offer new possibilities for creative expression and cutting edge, photorealistic visualizations, requiring minimum amount of time and effort.

Design and architecture, have embraced TTI generators with curiosity and both excitement and skepticism (Albaghajati et al., 2023). The allure of achieving photorealistic results in seconds, a task traditionally time-consuming, seems to intrigue designers to a great extent. As TTI usage is still evolving, lacking rigid protocols or standardized implementation methods, designers engage in experimentation to uncover the potential outcomes of interacting with these tools. Therefore, a lot of research interest emerges among this area in order to evaluate the prospects of it (Paananen et al., 2023), hoping that generative tools can enrich the design process by facilitating serendipitous discovery of ideas and nurturing an imaginative mindset.

5. Self-Knowledge through Iterative Explorations

As stated above, creating images with AI visualization tools takes place in very user-friendly interfaces and may seem quite straightforward initially. Users describe their desired image in natural language, specifying details like format, context, and subject matter. However, this process is not always as linear as expected. Achieving the desired outcome requires users to assess the output and refine the initial prompt until they discover the precise combination of words that conveys their vision. This transforms the process into a dialogue, an ongoing exchange between the user and the algorithm to attain the desired creative outcome. In the area of neural architecture, this technique allows shape to be interrogated through language (del Campo & Manninger, 2022).

This process introduces a dynamic interplay between the designer and the AI platforms, posing two critical questions. Firstly, there's the challenge of effectively conveying the designer's mental image through words. Secondly, it questions whether the AI accurately interprets and comprehends these verbal descriptions. The act of describing ideas with words transforms into a design process itself, fostering a continuous exchange. In this back-and-forth, the designer experiments with prompts, evaluates image results, and iterates on the process.

"Iterating" is an intrinsic part of design (Wynn & Eckert, 2017). Ideas and concepts find their form by molding through the laborious process of re-thinking, re-evaluating, re-sketching until one reaches the point where satisfaction and certain criteria are met. This way of operating is common in most designers. Through all these iterative steps, many variations of the project are being produced as part of the design and conceptualization process. This methodology of creation holds strong similarity with the way someone uses and explores ideas in AI generator platforms (Joyce, 2021). One can feed the same prompt into the algorithm and get different results. In advance to that, several platforms generate variations within the same prompt (Midjourney), providing the choice to the user to pick and re-iterate on the one that is deemed more preferable – somehow guiding them to their ideal outcome. AI's ability to quickly generate visualizations can play a key role when "iterating" on architectural design. The capacity to generate alternative design possibilities effortlessly is encouraging architects to explore new ideas and experiment with concepts and forms.

Crafting the 'perfect' prompt becomes an art of posing the right questions. It demands an understanding of effective communication with the AI, selecting keywords that precisely express intentions. Architect's engagement with different iterations reveals personal preferences, and self-knowledge emerges from recognizing design inclinations and biases. This interaction serves not only to refine the user's concepts but also to enrich and shape their design identity. The designer refines prompts to better convey their vision, creating a symbiotic relationship between human creativity and machine-generated insights. The focus extends beyond the final high-fidelity image; it encom-




Figure 1. "William Garner's visualizations on architectural concepts, produced via natural language descriptions using text-to-image generators (2022). The initial variations were produced using Midjourney and targeted editing took place using DALL-E 2. "

passes the evolving interaction with the tool, providing a new approach to self-reflection and concept articulation.

6. Workflow Transformation

In the field of architecture, the journey from conceptualization to construction involves navigating through various stages, each with distinct objectives. Architects lean on specified structured workflows to guide them through these phases (RIBA, 2021), with the technical aspects demanding a more rigid process, while the early stages allow for a flexible approach driven by creative thinking and intent.

The integration of digital design tools and photorealistic rendering into the design process has prompted architects to utilize visualizations for effective communication with clients, engineers, and other stakeholders. Achieving a photorealistic image entails a sequential process that essentially involves designing the building to a certain extent. While not a strict protocol, a common practice involves creating a brief, conceptualizing an idea, drafting initial sketches manually or digitally, progressing to resolving drawings and synthesizing the building, and finally constructing a 3D model as the foundation for photorealistic images.

However, a notable shift is occurring with the emergence of text-to-image algorithms powered by generative AI. These algorithms enable architects to provide a textual description of a building, generating remarkably realistic results in a short span of time, usually within seconds. This algorithmic approach appears to almost instantaneously design a structure, complete with materials, textures and atmospheric qualities, offering a glimpse into the potential

look and feel of the space. Therefore, the ability to introduce and experiment with early step visualizations prior to sketching, drawing and modeling any actual building can be seen as an opportunity to explore concepts and ideas without investing time and effort that would otherwise be prohibitive. This approach can be implemented into several stages of the design process including 1) General Study/Analysis with attention to storytelling and visual representations 2) Ideation/Synthesis with references and design inspirations and 3) Design development with multimedia and presentations. (Albaghajati et al., 2023)

The appeal of effortlessly acquiring impressive images might create a misconception that design is confined to the final visual representation alone. The risk is succumbing to cinematic effects without substantial design input, echoing a long-standing issue with the use of photorealism in design (Nastasi, 2016). However, portraying the integration of these tools as a threat (Leach, 2023) pushes us towards asking different kind of questions. Design is inherently multifaceted, and the potential of these tools to deliver exceptional results doesn't hint at an upcoming automation utopia. Instead, it encourages designers to view it as an opportunity—an open question challenging them to explore integration methods that enhance rather than undermine the design process.

One can certainly argue that text-to-image AI models and the features they offer can optimize production time in renderings using automation. The broader question is posing is looking into how- and if - it can be implemented into concept design stage to augment our capacity of creative thinking, rather than just facilitating photorealism.




Figure 2. “The shift of visualizations’ implementation during common architectural workflows at concept and schematic design stages.”

7. Creativity and Concept Making

As discussed above, one effective application of text-to-image AI generators is streamlining time-consuming tasks, particularly in the refinement of “rendering” and the creation of photorealistic visualizations, which were traditionally perceived as laborious and time-intensive tasks. More than that, the iterative nature of engagement with such tools showed the ability to enhance a designer’s self-reflective approach in design. But more than that, do these tools possess any potential to fundamentally reshape the way we approach design thinking? Rather than simply expediting established practices, can these tools inspire novel ideas and creative concepts that may not have emerged through conventional methods? This inquiry challenges us to explore the transformative possibilities that extend beyond time-saving measures and consider whether these tools contribute to a more profound evolution in our creative processes.

In architecture, referencing and quoting from the past to inform new designs has been a longstanding practice, notably embraced by post-modernist architects who creatively reinterpret historical architectural forms within their own work, challenging notions of originality and authorship. This approach serves as a form of meta-representation, reflecting upon the act of representation itself. When algorithms in text-to-image visualizations are trained to understand and interpret concepts, they essentially create a

representation of them. Subsequently, they translate these representations into visual images. This process constitutes a second-level representation, where the original representation is reconsidered or represented anew. This mode of operation, characterized by referencing, quoting, and critiquing existing architectural forms and ideas, aligns with what can be termed as content-aware metarepresentations (Moras, 2020). The aim of content-aware metarepresentations is to engage with architecture’s historical and theoretical heritage while exploring fresh interpretations and expressions within established frameworks.

Therefore, the main objective here is far from asking if AI can itself be creative. The big question is how can we get more creative by implementing such tools into our workflows, moving beyond the fear we shall be amputated by them. The scope is to pay attention to the ways we can utilize these tools to extend and expand the notions and ideas we already have, and see ourselves as augmented entities of creativity. Do such tools have the capacity to impute and ignite more creativity into our thinking or the way we use them only iterates and reflects on our very own human creativity? If the latter is the case, then one should consider such tools only as means for time efficiency and optimization in our creative production.

Exploring the alternative perspective, the proposed methodology draws inspiration from the creative practices of the surrealists, particularly the art of surrealist collage. Surrealist collage is a technique that enables artists to

swiftly combine and juxtapose pre-existing elements, be they words or images, to shape entirely new compositions. This seemingly irrational combination of disparate elements mirrors the construction of dreams, where unrelated elements converge to form peculiar narratives and scenes. (Cramer & Grant, 2020) The Surrealists viewed collage as a method to enact what they deemed the fundamental poetic activity of the unconscious mind: the fusion of diverse entities to generate something entirely novel. The Surrealist approach to collage manifested unconscious thought, serving as a conduit to authentic creativity. While these algorithms don't function precisely like collages, as they are trained to learn concepts via associating text to images (a process analogous to a child's learning), the objective is to leverage the concept of assembling irrational conceptual juxtapositions. These can serve as means of defamiliarisation and estrangement, offering the opportunity to observe the world differently than it is commonly understood or perceived (del Campo & Manninger, 2022). By experimenting with these juxtapositions on AI platforms, the aim is to explore how they can potentially deliver novel concepts and show genuine creative potential.

The example of the "elephant-fly" offers an introductory investigation, using a metaphorical language approach to explore unconventional ideas. It starts with portraying a large mammal as an ethereal being flying in the skies and refines the prompt to introduce a cross-breed animal species resulting from merging an elephant and a fly. While the initial idea is easy for humans to grasp, the latter concept demands a degree of unconventional thinking that may prove challenging even for the creative human mind. The juxtaposition that may perplex human creativity is effortlessly rendered by text-to-image AI, underscoring its capacity to generate imaginative constructs. The instances presented illustrate the algorithm's ability to seamlessly combine apparently unrelated ideas and blend elements from different taxonomies to craft imaginative visuals.

This approach doesn't exclusively depend on employing text-to-image AI; it can also be executed through traditional analog methods, a common practice in a designer's workflow. The age-old practice of "gathering references" when delving into a design problem, creating an inspiration board, and outlining design objectives is a well-established creative method. As life becomes more intricate, designers increasingly gather diverse and unconventional references, which can connote various elements like style, aesthetics, historical significance, material technology etc. For instance, let's consider a design concept integrating mycelium (a fungus with notable structure and mechanical properties which has been experimentally used as a construction material) either as a material or as a source of inspiration for structural formation. Suppose the same design concept aims to capture the qualities of a specific style, motion or expressiveness inspired by the Baroque movement. To illustrate, below are some reference images a creative individual might have collected. How might these references be translated into design ideas integrated into a spatial approach? Through an exploratory design process, a designer

can follow the traditional path using sketches, models, and the unlimited source of creative thinking.

Can text-to-image models facilitate a smoother transition from design references to spatial concepts, or even generate new "synthetic" references to expand the creative thinking of the designer? (In this context, "synthetic" denotes references not derived from real-life instances or existing works but rather created by algorithms.) The aim is to provide the algorithm with a prompt equivalent to the analog concept board's brief, seeking specific characteristics and qualities for the envisioned space. The generated images would then serve as an enriched concept board, showcasing initial digital spatial sketches – not a final photorealistic outcome. This approach harnesses creative visualization before embarking on any actual design steps, presenting potential scenarios that might escape the human mind.

Notably, there is an inherent loss of control in this process. The outcome from the algorithm remains unpredictable. Yet, it is within this uncertainty that the potential for creativity unfolds. Here, the designer willingly embraces an open-ended approach, allowing for the exploration of uncharted territory in the design process. Understanding the genesis of these tools allows for more meaningful integration into design workflows. Rather than being led by the technology, designers can leverage it to explore visions that might otherwise escape their purview. In essence, the medium becomes a conduit for design intention, capable of yielding both mundane and insightful results, hopefully surprising the designer with unforeseen possibilities.

8. Conclusions and Considerations

Can the generation of something "new" emerge from existing frameworks of creation? The approach to conceptual design in this manner poses a risk of disrupting the coherence and continuity of a design's conceptualization, potentially yielding a disjointed, fragmentary, and chaotic outcome. This risk is amplified by the absence of historic or social context in algorithmic creation and the blurred ethics associated with the massive datasets used to train these TTI algorithms. Critics argue that merging established aesthetic approaches with prompt-based images may lack the potential for true creativity and originality, as the process may be confined to replicating the work of other artists.

Considerations also surface regarding the interpretation of abstract concepts. Algorithms might represent the statistical average idea of "spatial expressiveness" based on their training data, potentially reflecting a specific cultural perspective (Naik & Nushi, 2023). Addressing such concerns necessitates a well-informed and sensitive approach from the research community and practitioners. It is crucial to establish ethical boundaries and understand the biases inherent in these creation workflows while also recognizing the power of collective imagination they offer to an open commons design community. Under certain conditions, these tools can foster creativity within the realm of distributed cognition and the synergy of collective imagination, laying the groundwork for innovative architectural expressions.




Figure 3. “Elephant flies VS Elephant fly: Evaluating the creative potential of text-to-image generators, by implementing surreal concept descriptions and playful language games. The images were generated using Stable Diffusion 2.1, a text-to-image model from StabilityAI.”




Figure 4. “Baroque Mycelium: Real life references collection for conceptualization of creative concepts. (Left: The Apotheosis of Hercules, baroque painting by François Lemoyne) (Right: Close up image of mycelium fungi structure.)”

Moreover, the concern that these tools merely reproduce old information without the capacity to pave the way for something new can be reconsidered. Alongside their unprecedented processing speed and graphic quality, these associative mechanisms hint at the potential for establish-

ing a new creative approach characterized by a sense of “hyper-dimensionality”. Since these algorithms are trained on datasets comprising myriad images from every historic era, they draw complex connections across a symbolic order that traverses time (Vickers & McDowell, 2021, p. 19). We




Figure 5. "Baroque Mycelium: Synthetic images produced with prompts including the keywords "baroque", "mycelium", "architecture", "expressiveness", "aggregations", "spatiality". The images were generated using Stable Diffusion 2.1, a text-to-image model from StabilityAI.

are witnessing the birth of a creative mechanism trained on the collective artistic knowledge of everything that has ever existed, capable of interacting simultaneously with the past, present, and future. Experimenting with these tools not only allows us to unlock new creative trajectories but also enables meaningful self-reflection, encouraging us to look back at cultural history and design in new ways. At this point, the idea of "collective imagination" is closely linked to "collective intelligence", which is no longer just a human phenomenon but is increasingly mediated and enhanced by technology (Morel, 2022). This technology does not just augment human intelligence but becomes an integral part of a new kind of collective intelligence and creativity that operates on a global scale, involving both humans and machines.

The last, very important aspect to be noted is the following: the algorithm created these images, but it didn't *design* them. The immediate question that arises is, who did? Did humans play a role by crafting the prompts that fueled the algorithm? Or should credit be given to the extensive dataset, the collection of images used to train the algorithm, incorporating the contributions of various creators? Another perspective suggests acknowledging the software engineers who dedicated effort and algorithmic expertise to create these sophisticated algorithms. Alternatively, one might attribute the outcome to the overarching technolog-

ical "infrastructure" and the engineers who meticulously designed electronic components with the necessary "memory" and "processing power". Is the algorithm itself deserving of recognition—an entity without sentience, yet capable of transcending material limitations to produce outcomes greater than the sum of its parts?

Regardless of the answers to these inquiries, this assertion posits that we confront an unbreakable interconnection of agents, each serving as an integral element within the whole, resulting in the emergence of a novel creative process. This horizon of unprecedented possibilities demands many considerations in terms of power, ethics and authorship, but its potential to provoke and propel the imagination in unforeseen ways already seems indisputable.

Funding

The research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) under the 3rd Call for HFRI PhD Fellowships (Fellowship Number: 5837).

Submitted: January 15, 2024 BST, Accepted: February 10, 2024 BST



This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CCBY-NC-4.0). View this license's legal deed at <https://creativecommons.org/licenses/by-nc/4.0> and legal code at <https://creativecommons.org/licenses/by-nc/4.0/legalcode> for more information.

References

- Albaghajati, Z. M., Bettaieb, D. M., & Malek, R. B. (2023). Exploring text-to-image application in architectural design: insights and implications. *Architecture, Structures and Construction*. <https://doi.org/10.1007/s44150-023-00103-x>
- Bernstein, P. (2022). *Machine Learning*. RIBA Publishing. <https://doi.org/10.4324/9781003297192>
- Chaillou, S. (2022). *Artificial Intelligence and Architecture: From Research to Practice*. Birkhäuser. <https://doi.org/10.1515/9783035624045>
- Cramer, C., & Grant, K. (2020, April 8). *Surrealist Techniques: Collage*. Smart History: The Center Dor Public Art and History. <https://smarthistory.org/surrealist-techniques-collage/>
- del Campo, M., & Manninger, S. (2022). Strange, But Familiar Enough: The Design Ecology of Neural Architecture. *AD Special Issue: Machine Hallucinations: Architecture and Artificial Intelligence*, 92(3), 38–45. <https://doi.org/10.1002/ad.2811>
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., ... Ozair, S. (2020). Generative Adversarial Networks. *Communications of the ACM*, 63(11), 139–144. <https://doi.org/10.1145/3422622>
- Joyce, S. C. (2021). AI as a collaborator in the early stage of the design. In I. As & P. Basu (Eds.), *The Routledge Companion to Artificial Intelligence in Architecture* (pp. 130–159). Routledge.
- Koh, J. Y., Fried, D., & Salakhutdinov, R. (2023). *Generating Images with Multimodal Language Models*.
- Leach, N. (2023, February 13). *AI is putting our jobs as architects unquestionably at risk*. Dezeen. <https://www.dezeen.com/2023/02/13/ai-architecture-jobs-risk-neil-leach-opinion/>
- Moras, A. (2020). Meta(re)presentations. *archiDOCT*, 14(7(2)), 14.
- Morel, P. (2022). Computational Intelligence: The Grid as a Post-Human Network. *archiDOCT*, 17.2(9 (2)).
- Naik, R., & Nushi, B. (2023). *Social Biases through the Text-to-Image Generation Lens*. <https://doi.org/10.1145/3600211.3604711>
- Nastasi, M. (2016). The City of Renderings. Photorealism, Spectacle and Abstraction in Contemporary Urban Landscapes. In R. A. Alcolea & J. T. Mingo (Eds.), *Inter - Photography and Architecture International Conference*. Universidad de Navarra.
- Oppenlaender, J. (2022). The Creativity of Text-to-Image Generation. In *Proceedings of the 25th International Academic Mindtrek Conference* (pp. 192–202). ACM Digital Library. <https://doi.org/10.1145/3569219.3569352>
- Paananen, V., Oppenlaender, J., & Visuri, A. (2023). *Using Text-to-Image Generation for Architectural Design Ideation*. <https://doi.org/10.1177/14780771231222783>
- RIBA. (2021, November 9). *RIBA Plan of Work*. <https://www.architecture.com/knowledge-and-resources/resources-landing-page/riba-plan-of-work>
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., ... Denton, E. (2022). *Photorealistic Text-to-Image Diffusion Models with Deep Language Unders*. <https://doi.org/10.1145/3528233.3530757>
- Shih, A., Sadigh, D., & Ermon, S. (2022). *Training and Inference on Any-Order Autoregressive Models the Right Way*.
- Vickers, B., & McDowell, K. (2021). *Atlas of Anomalous AI*. Ignota Books.
- Wynn, D. C., & Eckert, C. M. (2017). Perspectives on iteration in design and development. *Research in Engineering Design*, 28, 153–184. <https://doi.org/10.1007/s00163-016-0226-3>
- Zhang, C., Zhang, C., Zhang, M., & Kweon, I. S. (2023). *Text-to-image Diffusion Models in Generative AI: A Survey*.