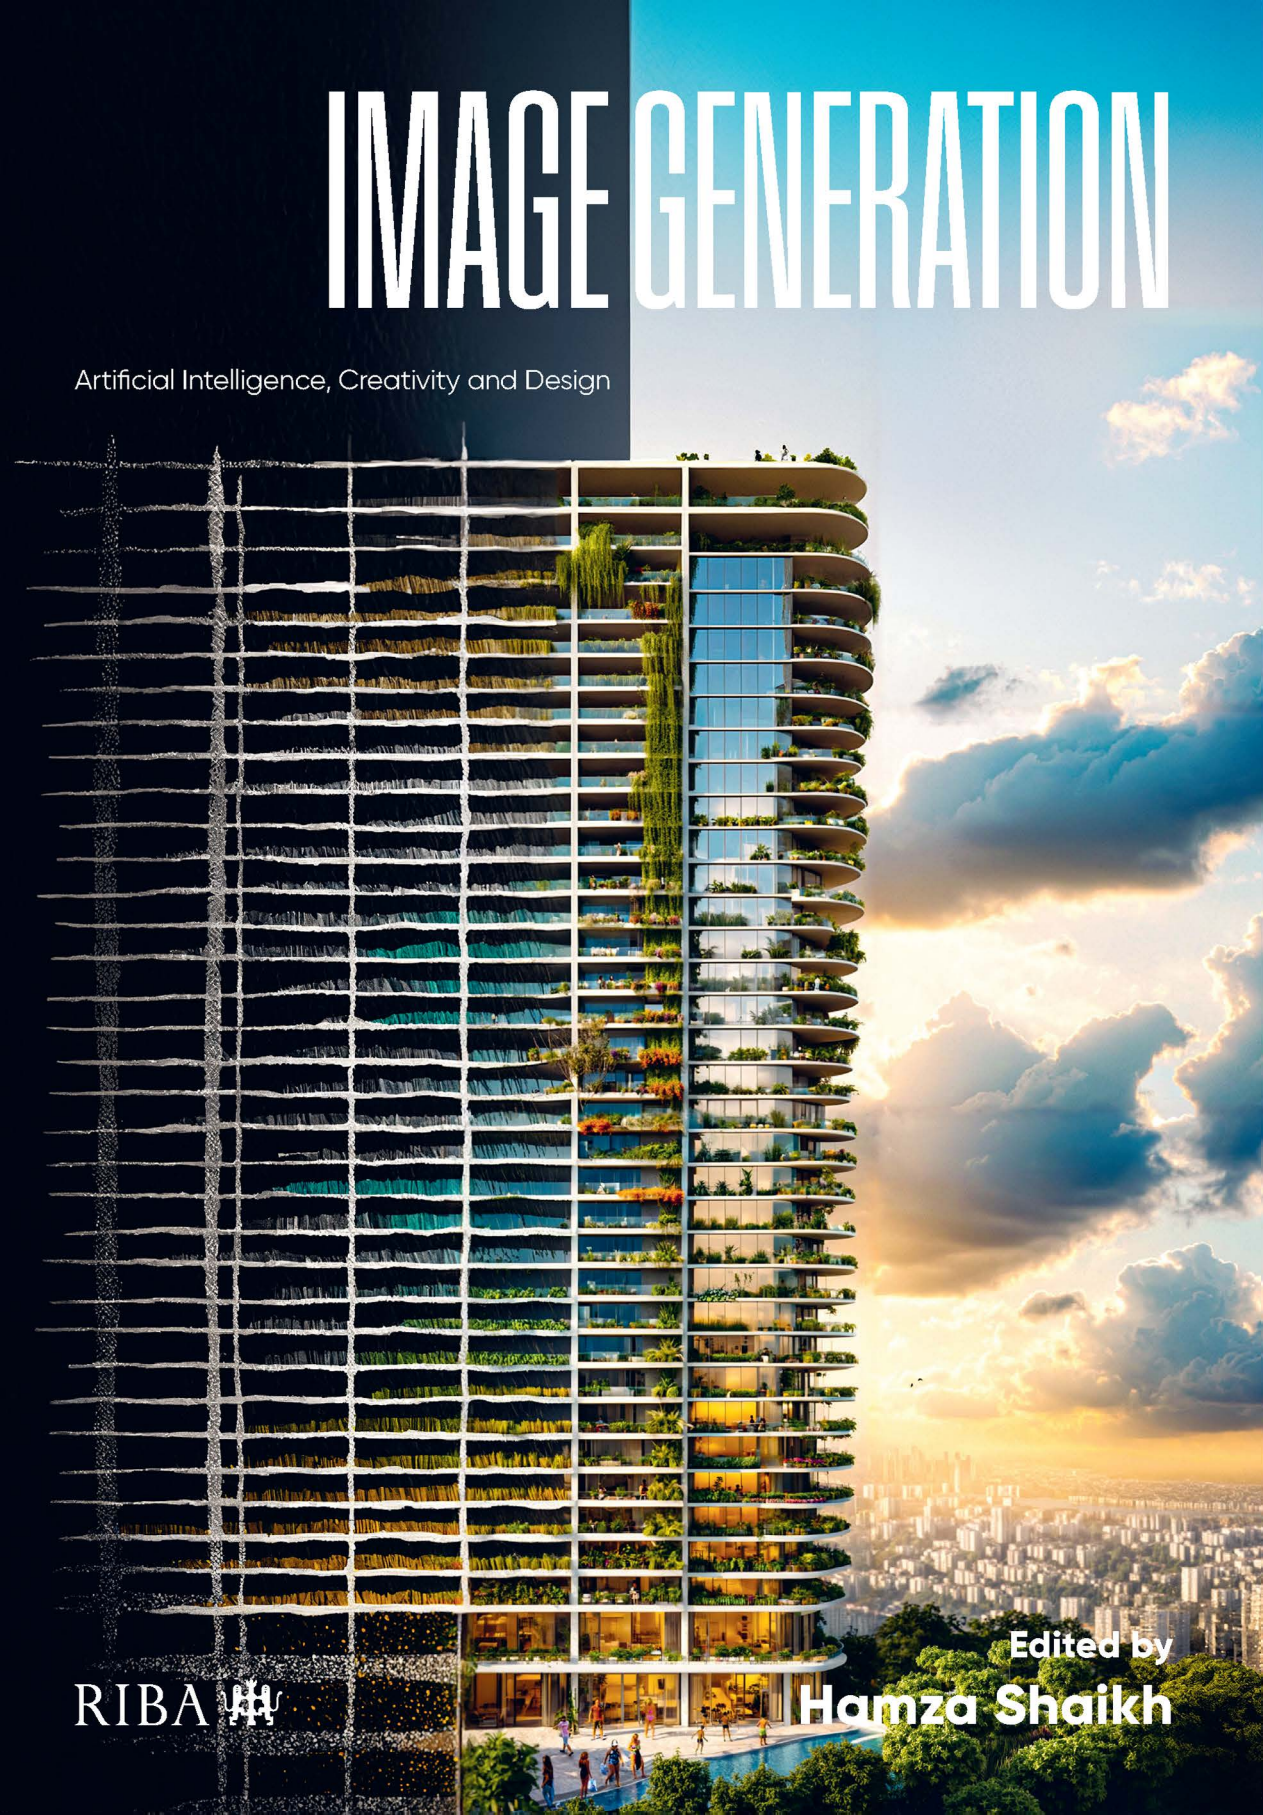
IMAGE GENERATION

Artificial Intelligence, Creativity and Design



Edited by

Hamza Shaikh

RIBA

IMAGE GENERATION

Artificial intelligence, creativity and design

Edited by
Hamza Shaikh

RIBA Publishing

Contents

About the Editor VI

Acknowledgements VII

INTRODUCTION IX
The AI Boom
Hamza Shaikh

T THEORY 01
A Picture is Worth a Thousand Words
Hamza Shaikh

W WORKFLOW 09
Generative 2D to 3D
TripoSR LRM + Blender + Rhino
Hamza Shaikh

W WORKFLOW 19
Laser Cut Prompts
Midjourney + Rhino + MDF Laser Cutting
Kat Stevens

T THEORY 31
Pixelstorms
Matias del Campo

W **WORKFLOW** 41
Real-Time Façade Iterations
ComfyUI + Grasshopper
Ömer Nuray

T **THEORY** 57
The Bias Dilemma
Hassan Ragab

W **WORKFLOW** 67
'I Prompted a Sofa, then Sold It!'
Midjourney + Rhino + Woodwork Upholstery
Rada Daleva

T **THEORY** 79
The Beginning of the AI
Revolution in Architecture
Tim Fu

W **WORKFLOW** 89
Breeding New Typologies
Midjourney + Kohya LoRA Training
Will Garner

T **THEORY** 101
The Effects of AI on Practice
Keir Regan-Alexander

Contents

W **WORKFLOW** 109
AI Form-Morphing
Stable Diffusion Automatic1111 Prompt Travel with XL LoRAs
Ismail Seleit

T **THEORY** 123
Architecture at 60fps
Carlos Bañón

W **WORKFLOW** 137
Sculpting with Inference
ChatGPT-4 + Processing 4.0 + ComfyUI
Carlos Bañón

W **WORKFLOW** 151
Real-Time AI-Enhanced Digital Collaging-to-3D
Krita + Stable Diffusion LCM Plugin + TripoSR v2.5 + Rhino
Hamza Shaikh

T **THEORY** 167
Beyond Three Dimensions
Thomas S. Sibley

W **WORKFLOW** 179
Robotic Mark-Making
ChatGPT-4 + Python + Grasshopper + KUKA|prc
Sara Codarin

T **THEORY** 191
Data is the Answer, but What is the Question?
Hamza Shaikh

T **THEORY** 199
AI and the Climate Challenge:
The Data Centre Paradox
Aleksandar Bajic & Hamza Shaikh

S **STORY** 209
Project Genesis: A Short Story on the
Future Roles of Architects in 2040
Hamza Shaikh

About the Contributors 218
Index 220
Image Credits 223

ABOUT THE EDITOR

Hamza Shaikh is a digital artist, architect and writer at the forefront of creative AI technology for design. As a Creative AI Lead at Gensler, one of the world's leading design firms, he develops human-centered AI solutions for the design process, driving innovation in concept design, storytelling and ideation on global projects. As part of Gensler's Design Technology Co-CEO HUB, he also leads training initiatives on generative AI, equipping thousands of architects and designers with the tools to shape the future of design.

In 2024, Hamza was ranked among the Top 5 in the Global Prompt Engineering Championship, solidifying his influence in AI-driven creativity. In 2023, he was recognised as a RIBAJ Rising Star for his thought leadership at the intersection of architecture, AI and social media. He also curated the acclaimed *Vanishing Points* exhibition at London's Roca Gallery in partnership with the historic Drawing Matter archive.

He is the author of *Drawing Attention: Architecture in the Age of Social Media*, which was RIBA Publishing's bestselling book in 2023, showcasing experimental drawing methods and exploring the role of digital platforms in contemporary architectural discourse.

Hamza holds First-Class Honours degrees in Architecture (BArch, MArch) and a postgraduate certification from the University of Cambridge. His work continues to shape the evolving relationship between AI, creativity and design.

ACKNOWLEDGEMENTS

Bringing this book to life has been both a challenge and a privilege. The past few years have been a whirlwind – navigating AI's rapid evolution, transitioning into a new professional role and preparing for a growing family. Yet, in the midst of it all, I've found immense fulfillment in exploring AI's creative potential – experimenting, developing workflows and testing ideas both at Gensler and in my own artistic pursuits.

Curiosity has always been my driving force, particularly at the intersections of art, technology and design. This book reflects that passion, shared by the contributors who, like me, are eager to push creative boundaries and embrace the unknown. I've been fortunate to explore AI playfully while applying those insights in a world-leading design firm. At Gensler, I've collaborated with experts and leaders on projects of immense scale, deepening my appreciation for AI's role in enhancing human creativity. The firm's entrepreneurial spirit and world-class digital innovation drive has allowed me to carve a unique path at the intersection of creativity, technology and design – an opportunity for which I am deeply grateful.

As with all my writing, I do this with students in mind. I remember the struggle of architecture school – balancing creative exploration with the demands of professional practice. The financial and emotional investment of those years fuels my sense of responsibility to uncover value for those entering the field. Though I am only 29 as I write this, fatherhood has heightened my awareness that we must offer the next generation more than we had. My hope is that students and professionals alike will turn to this book for clarity, direction and empowerment as they navigate AI's evolving role in design. I know that some of its content will eventually be outdated, but I believe it will serve as a marker of this pivotal moment – the first wave of 'hybrid AI' thinking in creativity and design. More than anything, I want this book to inspire others to carve their own paths by embracing creative experimentation for many years to come.

Finally, I owe my deepest gratitude to the strong, inspiring women in my life. To my mother, whose unwavering belief in me has given me the confidence to always strive for excellence and bold thinking. And to my wife, whose patience, support and sacrifices have allowed me to balance ambition with the joy of family. Their strength forms the foundation of everything I do.

It is with humility, gratitude, and hope that I share this work – knowing it reflects not just my journey, but the collective potential of all who are curious, courageous and committed to reimagining the future for a better world.

Hamza Shaikh
September 2025

VIII



INTRODUCTION:

The AI Boom

Hamza Shaikh

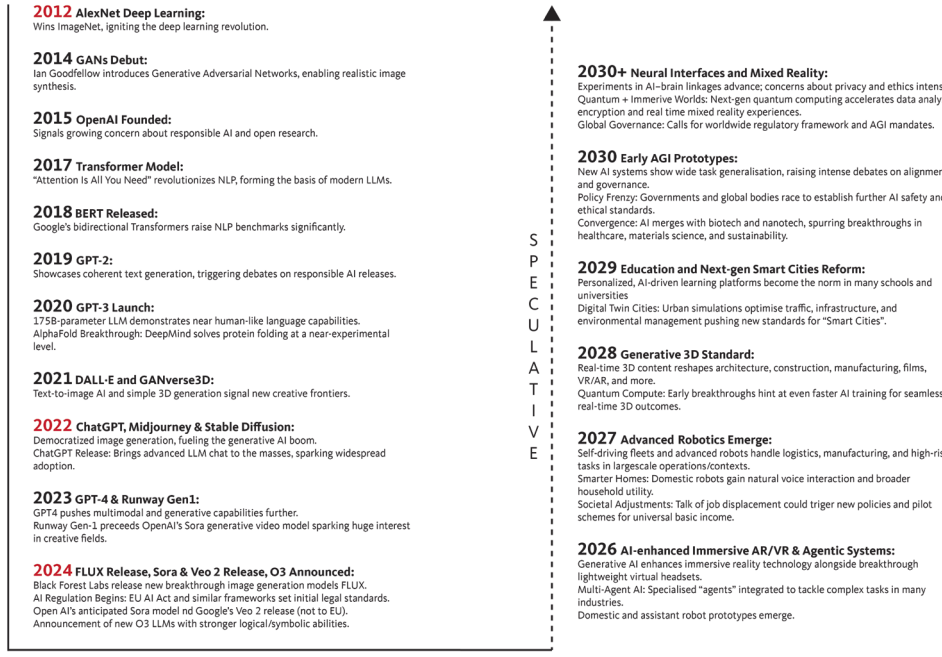
Today, we live in a world where attention is the ultimate currency. Platforms like Instagram and TikTok fuel a relentless stream of images, cultivating a culture of 'doom-scrolling' that feeds our escalating demand for ever-faster, immediate communication. It's no wonder, then, that the commercial AI boom was ushered in by a chatbot that could converse instantly and an app that could generate images at the speed of our imagination. Instant feedback and response aren't just hallmarks of our time – they are the fertile soil in which generative AI has firmly taken root.

In my previous book, *Drawing Attention: Architecture in the Age of Social Media*, I took a positive outlook on content creation, exploring how documenting ideation through drawings can fuel personal growth and opportunities for designers. The book also delved deep into diverse drawing styles and processes, featuring creatives who shared the technical and artistic intricacies behind their work. As I noted in that book, though one could easily write an entire volume on the downsides of social media (and many brilliant minds already have), my focus is on uncovering nuanced perspectives and opportunities for creativity. My hope is to guide people towards leveraging trends and disruptions as catalysts for growth, evolution and full-blown creative metamorphosis.

This book follows a similar intent to my previous one in providing theoretical ideas and technical insights into experimental design communication methods, however, we will dive into the heart of a new and profound transformation – the intersection of generative AI image generation with creativity and design. It's a shift with implications for rethinking our creative process at large and our relationship with technology. Over a period of a few years, the world has rapidly seen the evolution of the remarkable new opportunities generative AI technology offers, from commercial tools such as ChatGPT, DALL·E or Midjourney to advanced developments in computer-vision approximation, agentic systems and the signs of artificial general intelligence (AGI). At the time of writing this book (2024) the pace of development continues to grow, and AI is becoming significantly ingrained into our everyday life, work and discussions.

For countless people, OpenAI's release of ChatGPT in November 2022 served as a pivotal moment, sparking widespread public interest in AI and making generative technology widely accessible to the masses. While those in machine-learning fields had perhaps expected the rapid evolution of AI, the release of ChatGPT represented a transformative leap forward for the world, marking the beginning of a new era in technology and its social value. This was achieved through an extraordinary convergence of advancements: state-of-the-art research in neural networks, the monumental scaling of language models and NVIDIA's pioneering hardware optimisations, which offered the extraordinary computational power required to train and deploy such models. Together, these elements unleashed capabilities that few had anticipated seeing so soon, ushering in a wave of disruption where AI tools could engage, create and assist on an unprecedented level.

Arriving months before ChatGPT, image generation platform Midjourney offered powers previously unimaginable for creatives. Hosted, perhaps unexpectedly, on the social platform Discord, it could translate any combination of words into images within seconds, creating a new paradigm for accessible creative technology. Unlike DALL·E, which was initially limited to research circles, Midjourney fostered a collaborative community where users could easily generate and share visually striking art. Each successive version of its model brought more refined stylistic tools, heightening



A framing of a timeline on generative AI evolution and its speculative future, based on my writing merged with GPT-4 cross analysis and prediction.

anticipation with every new release. This social engagement, coupled with Midjourney's emphasis on artistic quality and its capabilities to control dimensions, style and creativity, introduced a rich world of image 'prompt engineering' – the nuanced, intentional crafting of inputs to achieve desired visual outcomes. These elements attracted creatives and led to viral growth on social media.

In only two years, Midjourney has reached over 20 million subscribers. As an early beta user, when it only had several thousand users, I was deeply engrossed by its potential. For me, witnessing its capacity to generate human-like drawings and intricate artwork marked an entirely new approach to creating mixed-media digital art with experimental drawing methods – which was my passion and background. As the app evolved to produce more photorealistic outcomes, I saw its potential to effortlessly traverse through mediums and design languages with nuanced prompting and varied inputs. I obsessed over the process of taking 2D sketches toward lifelike 3D representations – something my architect brain craved. Seeing a 'diffusion' process that expressed the progression of visual ideas and media was a new paradigm for design ideation and communication. At this point, and after seeing more people experiment and share their works on social media, it struck me that this technology could either be a risky tool for designers or a revolutionary one, depending on our approach. From then on, I focused on using it to create 'hybrid' works – those that merged my analogue artistic skills with the breadth and speed of AI's inferencing capabilities.

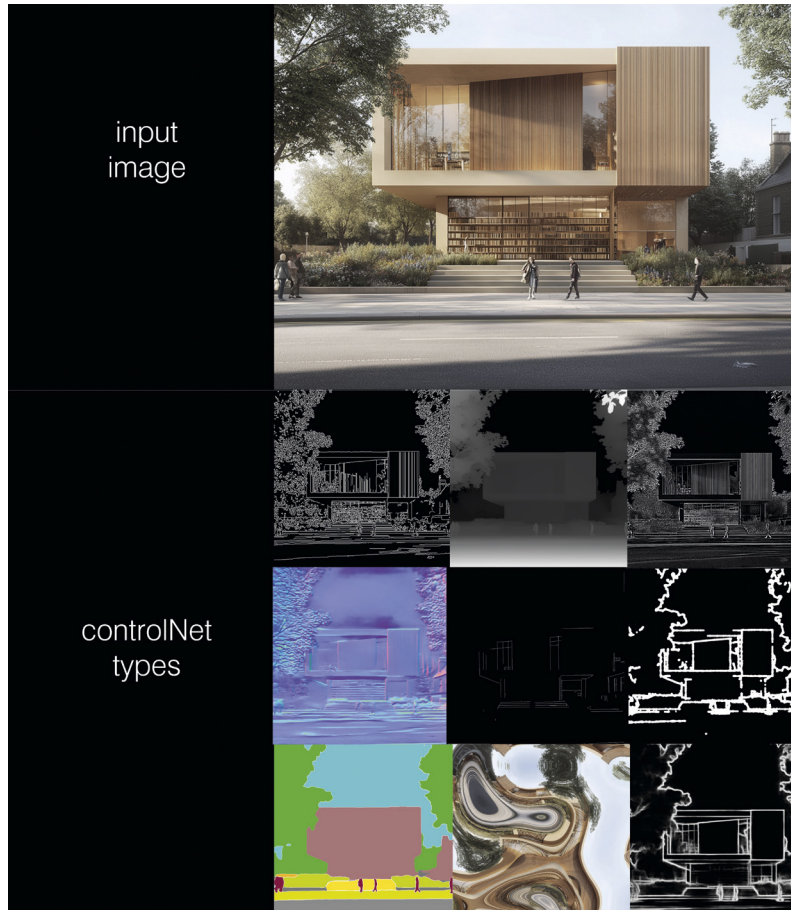


Two years' evolution of Midjourney through versions 1-6 using the same prompt.

While immersed in research and experimentation with generative AI, I was progressing through my career as an architect in London at the world-renowned design firm Gensler. This meant that my mind was deeply engaged in real-world design projects throughout the AI exploration. I was driven by a clear intent: to identify how these tools could enhance the value proposition for designers, architects and urban planners rather than diminish or threaten it. I eventually joined Gensler's global leadership team for design technology.

Our work involved piloting innovative workflows at one of the world's most innovative companies, granting me invaluable insights into the complexities of AI integration and the importance of data. It also underscored the importance of creating purpose-driven solutions that align with both human creativity and technological advancement. At Gensler, being a digital differentiator meant forging our own path for innovation with clear purpose and intent. Inspired by this, our responsibility to protect our data and to gain more advanced control, I looked to explore beyond commercial tools and platforms. To my delight, during this time, the Stable Diffusion repository was officially released.

With Stable Diffusion we unlocked a Pandora's box of ideation power and control. By experimenting with Stable Diffusion through its Automatic1111 and ComfyUI interfaces, we expanded our creative possibilities with a range of raw controls, each customisable and configurable to suit our needs. However, it was the introduction of an image translation algorithm called ControlNet that truly captivated the growing community of early adopters. This machine-learned function, integrated within the Stable Diffusion ecosystem, enabled the detection of image properties such as lines, depth, objects and subject matter through specialised feature extraction and encoding. Here, neural networks identify and organise these elements into structured data layers which then serve as guiding inputs, allowing Stable Diffusion to use our sketches or any chosen images as direct guidance for AI generation. Injected into the image creation process as a guiding mechanism, these dynamic data layers provide multi modal control. In the right hands, this technology yields incredibly powerful results, offering artists and designers the freedom to develop unique and inventive visual workflows that can blend with any human input of their choosing.



Stable Diffusion ControlNet types showing image translation previews.

The deliberate merging of human and AI capabilities forms a central theme in this book, encapsulated in what I refer to as the 'hybrid' approach. This approach advocates for the integration of nuanced generative AI functions with the irreplaceable qualities of human creativity, intuition and insight. While the hybrid philosophy is not entirely new, its relevance is growing across industries and professions where AI automation challenges human agency and accountability. Embracing a hybrid AI approach is not merely my recommendation for the future – it is an inevitability.

This book spotlights the first wave of pioneering talent, showcasing the hybrid approach in the architecture and design industries. The contributors, through theoretical stances and case studies present a diverse range of approaches to the hybrid process. Some focus on technical workflows and innovative design methods, while others present profound theoretical and philosophical stances – both equally significant in redefining the field. Despite the differing opinions shared, the featured contributions reveal how AI is not a replacement for human creativity but an extension of it – a complementary layer that enhances human expression, if used intentionally. By tailoring tools and workflows to align with individual creative visions, these innovators achieve outcomes that, while varying in technological sophistication, consistently excel in inventiveness and artistic flair. Together, they demonstrate the expansive potential of these tools to reshape not only creative practices but also how we think about design and its future.

A profound shift is underway: technology isn't just enhancing efficiency – it's creating. It's generating entirely new ideas, composing original text and crafting unique visuals. Where we once turned to books or scoured the internet for inspiration, we now wield tools that generate ideas both independently original and intimately connected to our own thinking. The power of AI in the creative process lies in its unprecedented capacity to expand our ideation and to do so collaboratively. With the right tools and mindset, we can reach solutions that previously felt unreachable and arrive at them in a fraction of the time. Yet, the role of human input remains indispensable. By taking charge of AI technology and bending it to what we do best, we can access creative outcomes like never before. Throughout this book, we uncover many examples of this careful balance: one that harmonises machine-driven possibilities with human input to shape meaningful, resonant outcomes.

However, it must be said that uncritical adoption risks steering us toward a future where design becomes superficial and driven by surface aesthetics or 'AI slop' lacking any genuine substance. Even more concerning for designers



Semi-flexible ideation - the iterative development of a sketch into varied options and mediums using a 'hybrid' AI process.

is the gung-ho adoption of automated commercial tools proliferating online, which could lead us to a future where creativity is increasingly commodified, and where agency and accountability become murky. My goal is to offer a constructive vision for AI's role in augmenting human creativity – a future where AI integration enhances our human value proposition and equips us with a profoundly stronger capacity to address some of the world's most pressing challenges.

This journey is not without its challenges, yet it brims with transformative possibilities. As we push into the frontier of AI in design, we're given a unique chance to reshape our practices and expand the limits of what's possible. The way forward is rooted in synergy – between humans and machines, between tradition and innovation. In the chapters ahead, you're invited to delve into the diverse insights of leaders and emerging talent in this field of hybrid AI creativity – to be inspired. Through the insights and experiences of these early pioneers, you'll gain a nuanced perspective on how you might thoughtfully weave AI into your own work, all while honouring the rigour of human craft. Whether you're an architect, designer or creative visionary, this book will hopefully inspire and equip you to navigate this evolving landscape with clarity, confidence and purpose.

XVI



A Picture is Worth a Thousand Words

Hamza Shaikh

DOI: 10.4324/9781003726500-1

In the age of hyper-connectivity, where every fleeting moment is immortalised in pixels, the power of imagery is undeniable. Yet, age-old aphorisms warn us against its deceit: 'Don't judge a book by its cover', 'Looks can be deceiving', and 'All that glitters is not gold'. These sayings underscore a universal truth – images are both revelatory and illusory. They entice, inform and manipulate, often simultaneously. The paradox of visual culture is that while we distrust its surface-level nature, we remain utterly captivated by its grasp.

The word 'image' originates from the Latin *imago*, meaning a likeness or portrait, derived from *imitari*, to imitate. This etymology hints at the duplicity of images: they represent reality yet are inherently fabricated. Sight – our most dominant sense – is as much a tool of discovery as it is a mechanism of deception. Historically, scientific discourse celebrated sight as a critical pathway to understanding, with the Renaissance elevating observation to the pinnacle of empirical inquiry and discovery. This reverence for the visual laid the foundation for a profound transformation in how knowledge was pursued and communicated. Fast forward to today, and entire industries thrive on the meticulous crafting of images – from the idealised perfection of social media feeds to the evocative allure of advertising campaigns. In this visually dominated landscape, imagery doesn't just inform, it influences, shaping our decisions, evoking emotions and crafting the narratives that define personal and collective identities.

But herein lies the tension: we criticise imagery for its superficiality while relying on it to navigate our world. It is not a mere coincidence; over 80% of human perception is shaped by sight, making it our primary channel for absorbing information. Neuroscientific research reveals that even those born blind experience visual-like processing through dreams and cognitive substitutes, underscoring the universality of visual thinking.[1] Vision's immediacy – its ability to deliver holistic impressions in milliseconds – sets it apart from sound or touch, which unfold over time. This efficiency has cemented its role as the backbone of creativity and design.

From the first flickering flames illuminating cave wall paintings to the pixel-perfect renderings of today emanating from our screens and virtual headsets, the evolution of creative practices has been inseparable from advancements in visual tools. Each leap – from charcoal sketches to immersive 3D worlds – reshapes the boundaries of human expression. Now, with generative AI, we stand at the brink of another revolution. This technology, capable of producing new images and narratives from text-based prompts, is redefining the essence of creativity itself. Generative AI is not an external entity but rather a tool that offers an extension to our creative capabilities, enhancing how designers think, iterate and execute by offering rapid and diverse avenues for exploration.

The Historical Weight of Visual Culture

For centuries, visuals have served as humanity's universal language, transcending cultural and linguistic barriers. The Renaissance, with its emphasis on linear perspective and realism, shifted art and science into a visually empirical realm.[2] This era wasn't just about artistic beauty but about redefining how humans understood their place in the cosmos. Galileo's telescopic drawings of the moon and Da Vinci's anatomical sketches exemplified the power of visuals to bridge observation and understanding.[3]



Depiction of cave painting and early carving/drawing tools.



Depiction of Narcissus looking at his reflection on an iPhone above the water.

This tradition carried into the industrial revolution, where technical drawings and blueprints became essential to engineering and architecture. The capacity to translate an idea into a visual medium – accurate, precise and replicable – underpinned the rapid advancements of the era. As photography emerged, it further democratised imagery, capturing reality in ways that were once limited to the artist's brush.

The 20th century saw the explosion of advertising, cinema and television, embedding imagery deeper into daily life. Visuals weren't just representations, they became tools of persuasion, emotion and narrative. The rise of social media in the 21st century transformed images into powerful commodities, central to both personal and collective identity. Instagram filters and YouTube thumbnails shape self-perception and engagement, highlighting how visuals have become tools of influence in every sphere of life.[4] Today, individuals are increasingly empowered to take control of image creation, crafting narratives that define their roles in business and social contexts.

This shift taps into humanity's deep-seated desire to be seen and to see others, fostering meaningful connections in a hyper-connected world. With the ability to reach global audiences overnight, visual storytelling has become essential for self-expression and influence. This democratisation of image-making places the most advanced methods of storytelling, branding and visualisation directly in the hands of creators.

Generative AI and the Iterative Revolution

Generative AI models like DALL·E, Midjourney, Stable Diffusion and Flux have amplified this iterative image-making potential exponentially.[5] These tools can produce countless variations of a concept in moments, breaking down the barriers of time and manual

effort. Designers can now explore pathways previously unimaginable, pushing their creativity to unprecedented heights. The saying ‘a picture is worth a thousand words’ has never been more relevant. With today’s technology, an entire chapter from a book can be transformed into an evocative image, or a single-worded prompt can generate thousands of varied visual interpretations. This capability underscores the expansive potential of AI tools in translating ideas into visuals at unprecedented levels.

The true potential of AI lies not just in generating a multitude of images but in its ability to empower us to craft compelling narratives, refine intricate ideas and shape visionary concepts. Mastery comes from iteration – the deliberate transformation of raw outputs into meaningful creations that resonate deeply. While basic prompt-writing skills can yield visually appealing results, true innovators push beyond the surface. They leverage precise controls, seamlessly merging traditional design principles with advanced digital tools creating ‘hybrid’ outcomes. They cultivate an experimental mindset, essential for driving breakthroughs in creativity and design, and recognise the importance of carefully curated data to train and refine models, ensuring thoughtful, intentional outcomes at every step. These principles not only unlock the full power of generative AI but also form the foundation of this book and the contributions shared.



A ‘hybrid’ process showing a hand sketch converting into a 3D-rendered image.

AI Inclusivity and Access

Beyond empowering individual creativity, AI profoundly enhances collective collaboration and accessibility. Text serves as a bridge between vision and execution, allowing diverse voices to contribute meaningfully to design ideation. This opens extraordinary possibilities for collaboration, where every team member can actively shape a shared vision. By brainstorming with words and seamlessly translating ideas into images, teams function like a creative hive mind, exchanging and refining concepts in real time. Generative AI's instant feedback fosters inclusivity, integrating perspectives from all participants, regardless of technical expertise.

This dynamic not only drives innovation but also reinforces collective ownership of the final design, resulting in more cohesive and impactful outcomes. The collaborative nature of AI extends its benefits beyond individual contributions, unlocking tools and insights that fuel collective innovation. For visually impaired creators, AI bridges accessibility gaps by converting textual descriptions into vivid imagery, enabling participation on an equal footing with seamless text-to-image-to-text possibilities.

As we will uncover, advanced generative AI models and platforms can also provide creators with unparalleled precision and flexibility, enabling fine-tuning of image models with bespoke data sets, live sketch conversion into multimedia outcomes and even the conversion of 2D images into 3D designs hinting at exciting future possibilities. Innovations such as segmentation-based ControlNet and open-source resources from industry leaders like Meta and Google continue to expand creative possibilities, democratising access to these powerful tools. Even without coding expertise, designers can now adapt AI pipelines to produce highly customised processes, aligning their craft with their unique creative visions and AI requirements.

The Future Beyond 2D

The trajectory of design tools points toward a future where AI-generated outputs transcend static, two-dimensional representations and redefine how we interact with robotics-enabled construction and virtual environments. Immersive technologies such as augmented and virtual reality are already being integrated into design workflows. Platforms like Unity, Unreal Engine, and NVIDIA Omniverse are at the forefront of this transformation, enabling the creation of digital twins[6] – highly detailed virtual replicas of physical spaces that leverage AI to offer an interactive simulation of real-world conditions such as climate, lighting, acoustics and spatial dynamics.[7]

Imagine walking through a digitally constructed environment, where natural light is simulated as it streams through windows and sound dynamics reverberate across walls in real time. This potential enables architects and designers to test and refine their ideas with unparalleled accuracy before committing to physical resources. These data-driven 3D and virtual environments significantly reduce financial risks, prevent costly design errors and promote a more sustainable approach to construction by optimising resource usage and minimising waste. This is particularly crucial given the substantial energy demands associated with technology production and maintenance, especially in AI systems. Addressing these challenges through smarter design processes ensures that technological advancements align with global sustainability goals. Later chapters will delve deeper into how these tools can mitigate environmental impacts and balance the energy footprint of advanced technologies.



Depiction of a future construction site with automated, generative 3D-printing systems.

This transformation goes beyond visualisation. AI-driven robotics are gradually reshaping the construction industry, with automated processes such as bricklaying, concrete pouring and large-scale 3D printing becoming more prevalent.[8] These advancements, integrated with data-driven design tools, enable streamlined workflows where architects and engineers can oversee projects from conception to execution with increased precision. This integration optimises resource allocation, reduces material waste and enhances efficiency, aligning construction practices with sustainability goals.

As we enter a future shaped by AI-powered technologies, construction is poised for significant shifts. Data insights and automation will not only improve project accuracy but also address challenges such as labour shortages and environmental impact. These changes signal a broader transformation in how buildings are designed, constructed, and maintained, emphasising resilience and sustainability to meet the demands of modern infrastructure.

Adapting to the Future

The rapid pace of AI development suggests that these tools will increasingly shift from assisting in creative processes to reshaping industries in fundamental ways. In construction, as mentioned previously, automation through robotics and machine learning is streamlining labour-intensive tasks such as bricklaying, concrete pouring and structural detailing. While these advancements reduce risks and improve efficiency in areas where human involvement may not be critical, they also free up human capacity for more creative and sustainable pursuits.

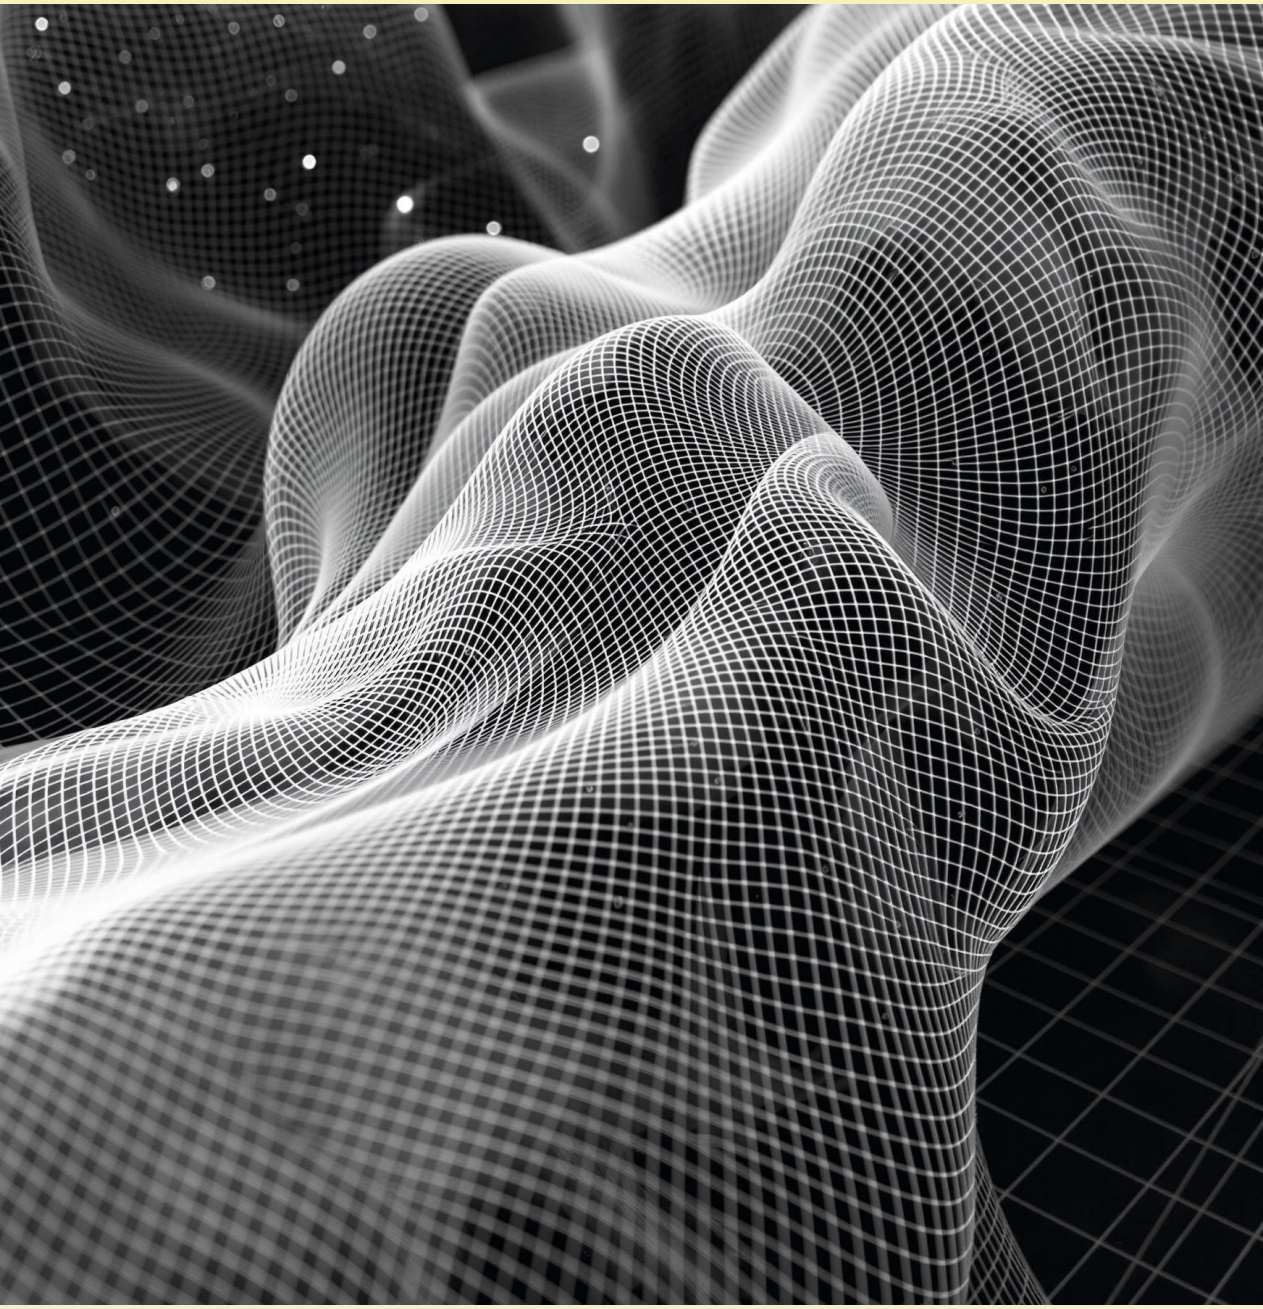
Despite the emphasis on automation, its most appropriate applications are in areas where it enhances efficiency and minimises risks, particularly in routine technical tasks where human error and oversight are frequent. For example, detailed technical drawings and structural modelling in architecture and engineering are increasingly optimised by data-driven AI algorithms, making these tasks well-suited for automation. However, this shift also creates an opportunity for professionals to redefine their roles in light of these advancements.

Architectural technologists and structural engineers, for instance, might evolve into technical data specialists, focusing on refining and managing generative 3D structural models.[9] This adaptation ensures their expertise remains integral to creative and technical workflows while leveraging AI to improve precision and streamline processes. Automation in these areas does not replace human involvement but rather reallocates it, allowing professionals to prioritise innovative and sustainable aspects of their work. Human monitoring and query will remain critical to ensure that automated systems operate ethically and align with creative and practical objectives.

This shift enables architects and engineers to focus on innovative, sustainable and creative aspects of their work while allowing automation to handle repetitive processes. By improving precision and efficiency in areas like construction and technical detailing, automation offers a more reliable framework for implementation. These changes raise an intriguing question: Are we entering a new creative renaissance, where technology empowers us to focus more deeply than ever before on imaginative problem-solving, strategy and sustainability? In any scenario one might envision, two things remain clear: generative AI is an inevitable part of the future, and how human ingenuity adapts to working alongside AI will fundamentally shape our roles and contributions in this evolving hybrid landscape.

Notes

1. G. Dormal & O. Collignon, 'Functional selectivity in sensory-deprived cortices', *Journal of Neurophysiology*, 105(6) (2011), [doi: 10.1152/jn.00864.2010], pp. 2627–2630.
2. P. Burke, *Eyewitnessing: The Uses of Images as Historical Evidence* (Ithaca, NY, 2001).
3. M. Clayton & R. Philo, *Leonardo da Vinci: Anatomist* (London, 2001) and ed. H. Bredekamp, 'Galileo's Thinking Hand: Mannerism, Anti-Mannerism and the Virtue of Drawing in the Foundation of Early Modern Science', in *Galileo's O* (Berlin, 2014), pp. 15–42.
4. C. Montag & S. Diefenbach, 'Towards Homo Digitalis: Important Research Issues for Psychology and the Neurosciences at the Dawn of the Internet of Things and the Digital Society', *Sustainability*, 10(2) (2018), [doi: 10.3390/su10020415], p. 415.
5. A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. & I. Sutskever, 'DALL·E: Creating Images from Text', OpenAI Blog (2021), https://openai.com/dall-e, accessed 27 Feb 2025.
6. Unity Technologies. (n.d.). Unity Real-Time Development Platform, https://unity.com, Epic Games. (n.d.). Unreal Engine: Powerful Real-Time 3D Creation Platform, https://www.unrealengine.com, NVIDIA. (n.d.). NVIDIA Omniverse Platform, https://www.nvidia.com/en-us/omniverse, all accessed 27 February 2025.
7. M. Grieves & J. Vickers, 'Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems', in F. Kahlen, S. Flumerfelt & A. Alves (eds.), 'Transdisciplinary Perspectives on Complex Systems', *Springer*, (2018), [doi: 10.1007/978-3-319-38756-7_4], pp. 85–113.
8. S. Lim, R. A. Buswell, T. T. Le, S. A. Austin, A. G. F. Gibb & T. Thorpe, 'Developments in construction-scale additive manufacturing processes', *Automation in Construction*, 21 (2012), [doi: 10.1016/j.autcon.2011.06.010], pp. 262–268.
9. Grieves & Vickers, 'Digital Twin'.



Generative 2D to 3D

TripoSR LRM + Blender + Rhino

Hamza Shaikh

DOI: 10.4324/9781003726500-2

JOURNEY

Generative 2D technology is rapidly evolving into more advanced generative 3D capabilities. Before the AI boom, existing solutions such as photogrammetry and videogrammetry were commonly used to create NeRF point clouds or Gaussian splatting models. However, recent breakthroughs in machine learning have introduced large reconstruction models (LRMs), which utilise diffusion-based predictive technology to synthesise results from photogrammetric scans. These models leverage perceived depth maps to stitch together 2D images into fully 3D-modelled meshes, offering a significant leap in automation and precision.

One notable example is TripoSR, an open-source LRM released by Stability AI. Trained on the Objaverse dataset, TripoSR is currently limited to generating basic 3D models with a 'miniature toy-level' detail and resolution. Despite these limitations, its potential is evident, as demonstrated through workflows that extend its capabilities. For instance, using TripoSR to transform a 2D-generated sofa design into a 3D model showcases the model's versatility in product design.

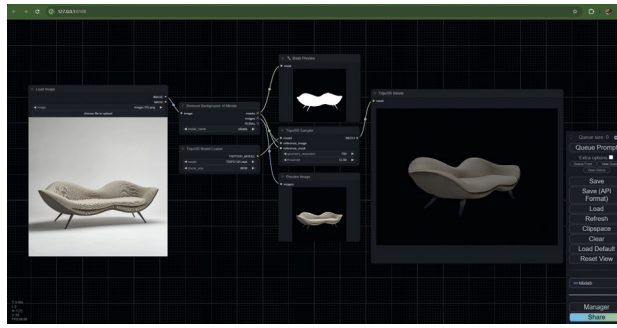
By adopting a hybrid approach – combining TripoSR with 3D design software like Blender and Rhino – I demonstrate how to quickly generate, clean and refine 3D assets. This workflow highlights how AI-powered tools can enhance product and interior design processes, enabling rapid prototyping and customisation while integrating seamlessly into creative pipelines.

WORKFLOW

STEP 1

To start, install and run ComfyUI. Once it's set up, configure the workflow using Mixlab nodes for smart background removal and TripoSR for image-to-3D conversion. This model works best with simple forms, such as a toy rubber duck or an action figure with basic details, but for a more interior/product design relevant task you can use a simple chair or sofa. I generated the 2D image of the sofa using Midjourney and ensured it was a recognisable object with no funny angles.

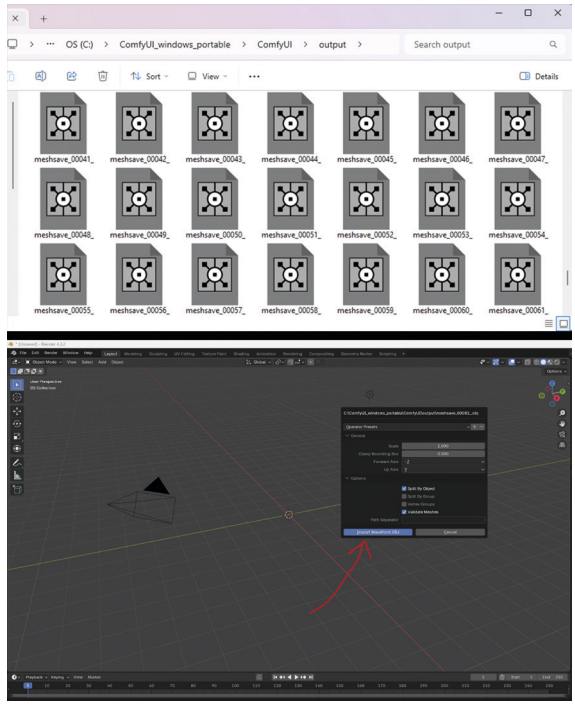
This setup was tested with 16GB of video random-access memory (VRAM) on a NVIDIA RTX 4070 graphics card, taking approximately 20 seconds to generate the 3D model. Adjust node settings as needed to match your hardware.



Set up ComfyUI nodes for img-to-3D.

STEP 2

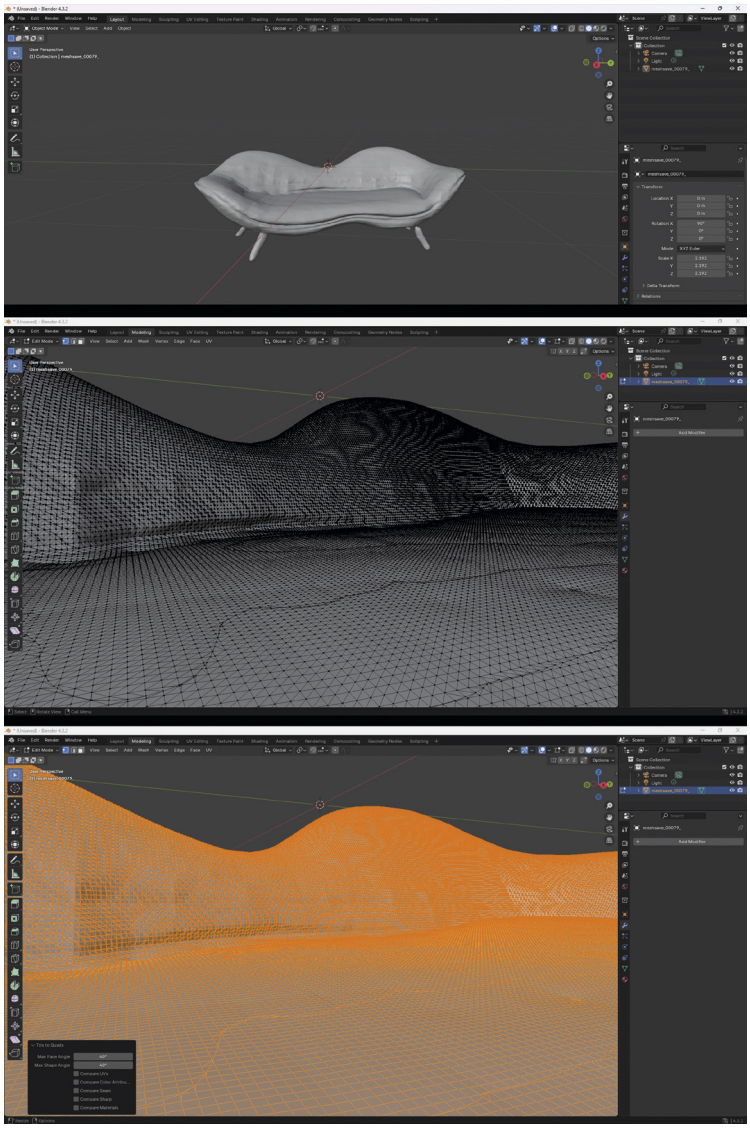
Locate your outputs folder from ComfyUI, typically found at *C:\ComfyUI_windows_portable\ComfyUI\output*, where files are often saved in Wavefront OBJ (.obj) format. Import or drag and drop the .obj file directly into Blender.



Import 3D output into Blender.

STEP 3

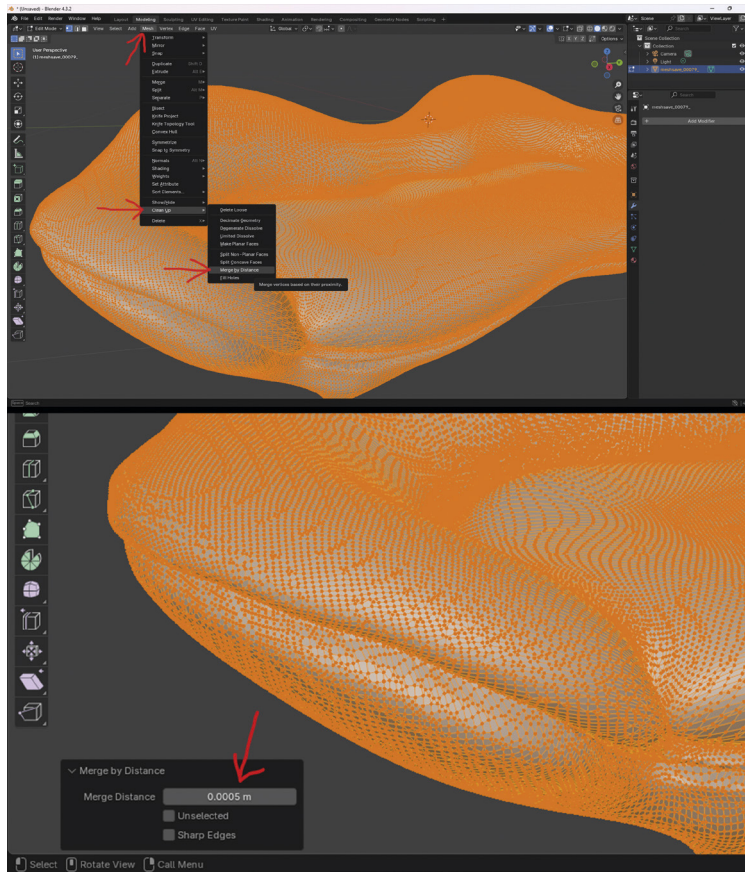
Once your 3D model is imported, it may appear messy due to the LRM image translation process. Begin by assessing the model to identify areas that need cleaning. Before making any manual smoothing or sculpting adjustments, convert the triangulated meshes into quads by entering ‘Edit’ mode, pressing ‘A’ to select all, and using the shortcut ‘alt + J’. This simplifies the geometry, making further cleaning and adjustments more efficient.



Clean model - turn triangulated mesh into quad mesh.

STEP 4

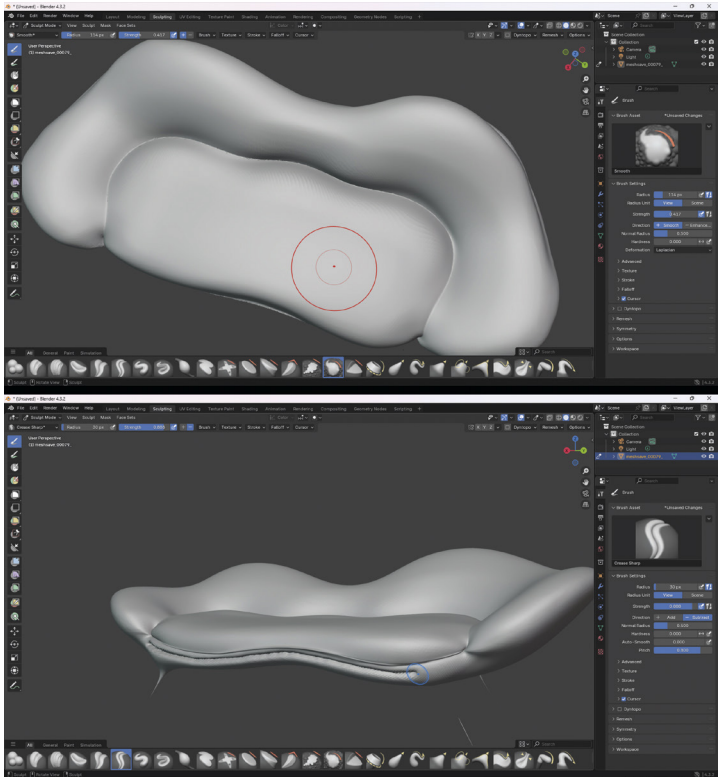
Continue cleaning the mesh by using the ‘Merge by Distance’ command. Navigate to the ‘Mesh’ menu, then ‘Clean Up’ and select ‘Merge by Distance’. In the parameter control box that appears, adjust the distance value; increasing it from the default 0.00001 to 0.00050 can be effective. Avoid setting it too high, as this may overly distort your model, but do what works best for you.



Clean model - merge by distance to reduce mesh complexity.

STEP 5

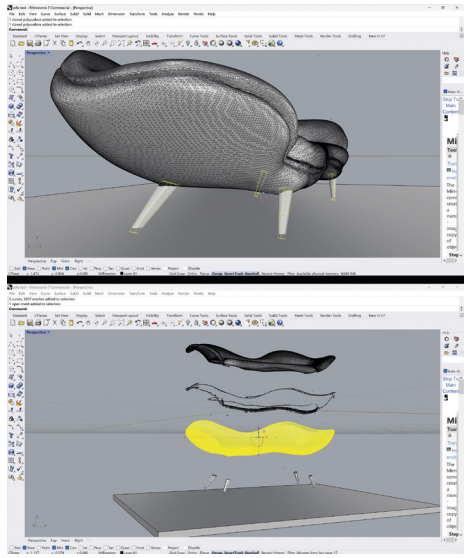
Now, refine your 3D model using Blender's sculpting tools. The 'Smooth' tool is useful for removing imperfections, while the 'Crease' tool can add detail. If excessive smoothing erodes the model too much, follow up with the 'Inflate' tool to restore volume. Ensure all parts are treated uniformly and symmetrically if needed. Once satisfied, export the model in .obj file format again.



Manually sculpt to refine as needed.

STEP 6

Import your .obj file into Rhino for further editing. For example, you can rebuild elements such as sofa legs and adjust proportions as needed. Separate the furniture into its constituent parts to assign distinct materials or refine them further for fabrication purposes. Use the 'Weld' command to merge vertices where needed, then 'Explode' to divide the mesh into manageable sections for easier manipulation.



Import into Rhino and edit further.

STEP 7

Here, you can see the starting image and the final result, where the image has been converted into 3D using a hybrid design process. The furniture can now have materials applied, rearranged and customised. This can then be rendered using visualisation software or image generation software with ControlNet, as shown in the example above (Automatic1111 + ControlNet).

Additionally, you can build a full furniture collection, which can be used to populate an interior model for visualisation or fabrication.



Set materials and scene and make variations of your design in 3D.

TIPS

Balance is everything	The success of any software lies in the tension between its power and usability. But more fundamental than that is who it was made for - sometimes, complexity isn't a flaw but a necessity.
Track the trajectory	AI is evolving at an exponential pace, but without understanding its origins and the evolution of tools and workflows, you'll always be playing catch-up.
Technology is more than a tool	It's a philosophy, a culture, and a reflection of our desire to extend human capability. What we create ultimately shapes who we become.

What does the future role of the architect or designer look like?

We are all becoming developers. The barriers to building digital AI tools have never been lower, shifting tech development from an exclusive skill set to an open frontier. What happens when anyone can create anything? A world where every idea can be coded into existence reshapes not just industries, but imagination itself. The power to build is no longer reserved for the few - it belongs to anyone with curiosity, a problem to solve and access to AI. This changes the game entirely. But with limitless creation comes a deeper challenge: discernment. The question is no longer just who can build, but what should be built? As AI removes friction from development, the true differentiator will no longer be technical ability, but vision - the ability to create meaning, not just functionality.



Laser Cut Prompts

Midjourney + Rhino + MDF Laser Cutting

Kat Stevens

DOI: 10.4324/9781003726500-3

JOURNEY

While studying architecture at the University of Westminster, I focused on reimagining the importance of how we might depict 'wellness' using generative AI as a tool. I was guided by my course tutor to push the boundaries of artistic explorations, and discover new ways to combine generative AI with common practices and methodologies taught in design school. My workflow combined hand drawing, Rhino 3D modelling/computer-aided design (CAD) and Midjourney image generation.

My step-by-step workflow showcases how these methods can work together in a single-concept design project. I discovered that generative AI imagery can be used as a tool to generate inspirational imagery with surprisingly accurate mathematical geometry. This process shows how I take these generations and try to refine them, eventually translating the imagery to create physical models.

WORKFLOW

STEP 1

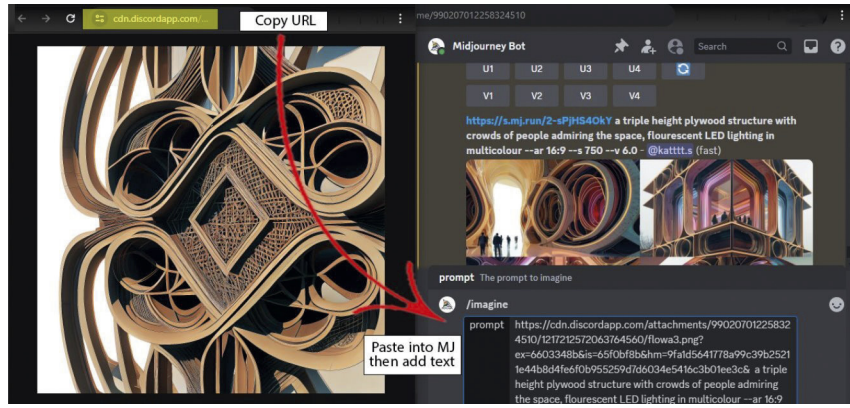
Generating the Midjourney image

Begin prompting in Midjourney using the concept of wellness as the underlying inspiration. Start with a basic prompt; you do not need extensive or highly technical prompts to achieve good results. Simply describe your idea in simple informal language as a great starting point. This process naturally becomes iterative once you see the resulting images from your prompts.

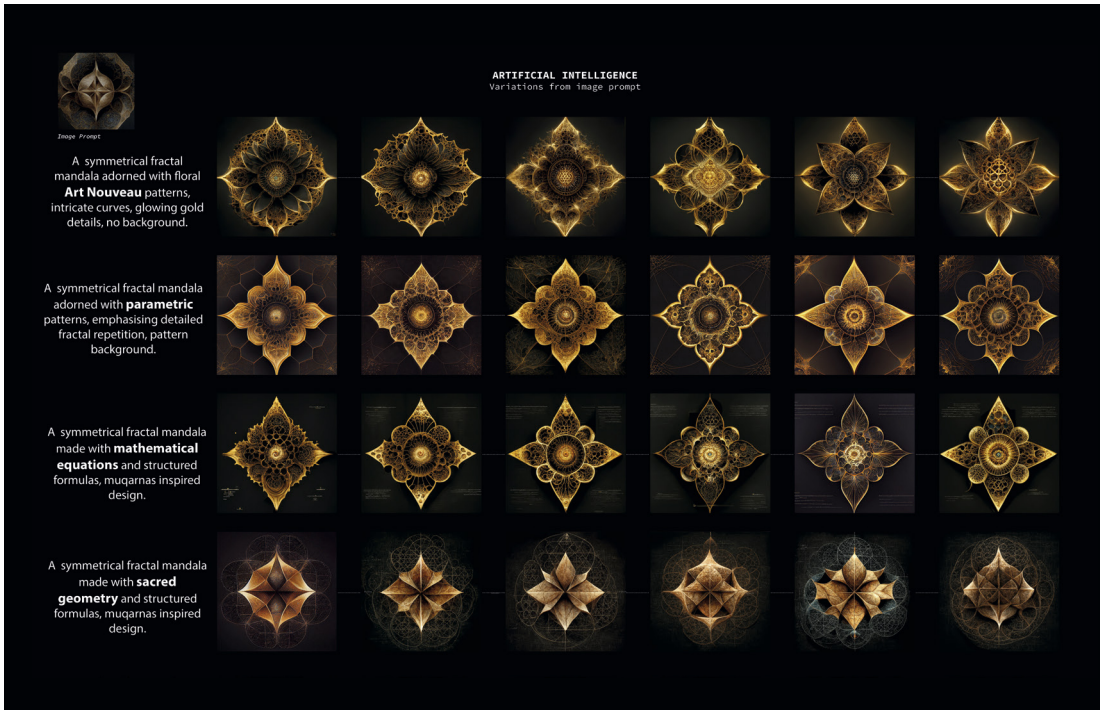
After getting satisfactory results, look for a reference image to guide and stylise the prompt. Think of the reference image as injecting the DNA of any image into your prompt using a unique URL. To insert an existing image into your Midjourney UI, drag and drop an image onto the text box, hit 'enter', open the image in a separate tab, copy the URL and paste it back into the text box after typing.



Enter a simple prompt.



Inserting an image as reference and the resulting outcome with prompt.



Continued prompt iterations using the image as a reference.

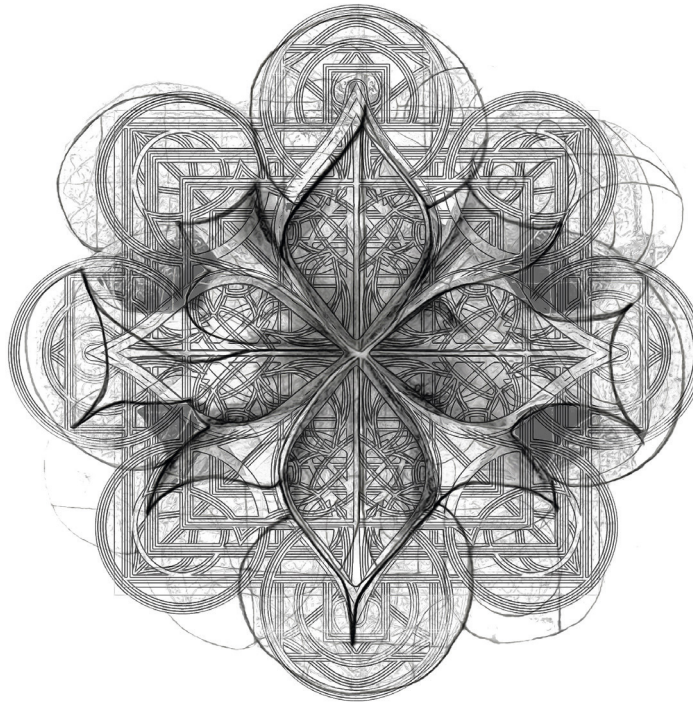
STEP 2

Midjourney to computer modelling

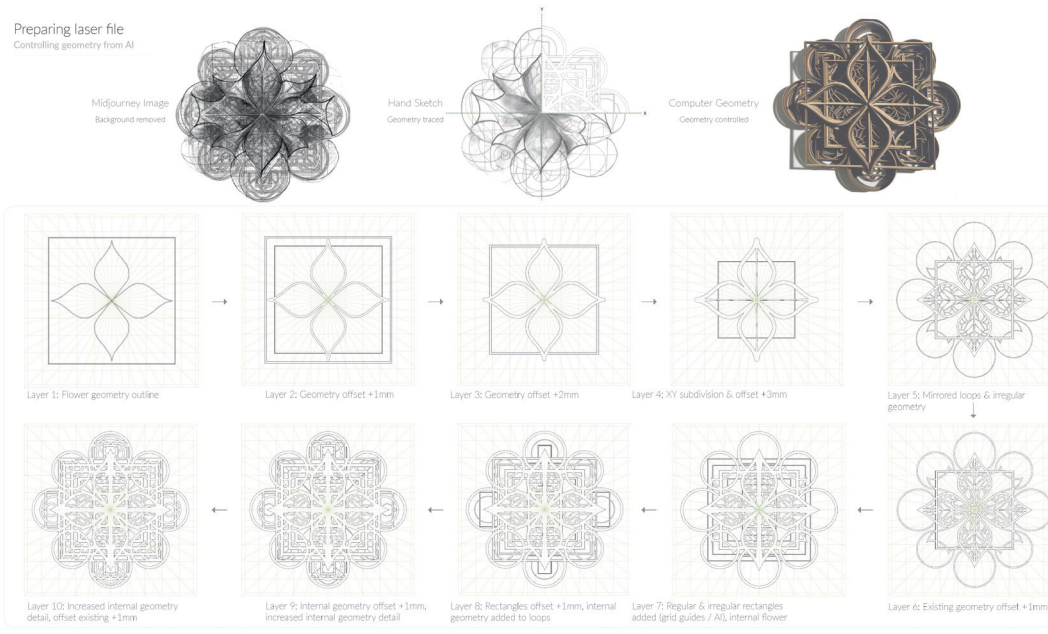
After extensive generation and curation, select the favourite outcome to begin exploring methods of translating geometry from the image into CAD format.

The key to separating the geometry into individual layers is to understand how to break down the geometry, which involves an interpretive and subjective approach. Start by sketching and re-sketching over the image by hand until the main geometric components are clearly understood. Then, assess how to separate these components into layers.

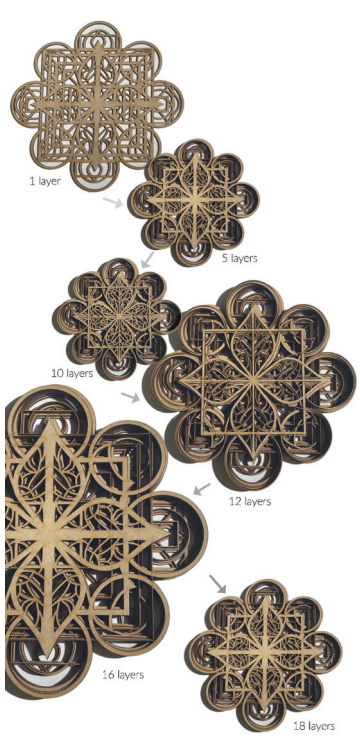
Import the resultant 2D sketch into Rhino. Using the ‘Curve’ tool, begin tracing over the image with your mouse, organizing the CAD linework into layers. Once complete, laser cut the layers onto individual plywood sheets. The result will be a handcrafted, laser cut layered model that expresses the AI-generated image as a motif.

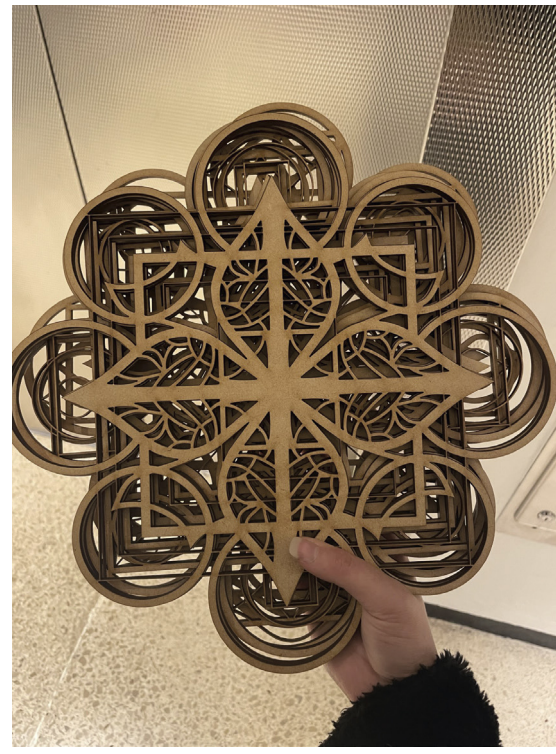


Trace and retrace over your AI image until you understand its geometry.



Break up your AI image into layers.

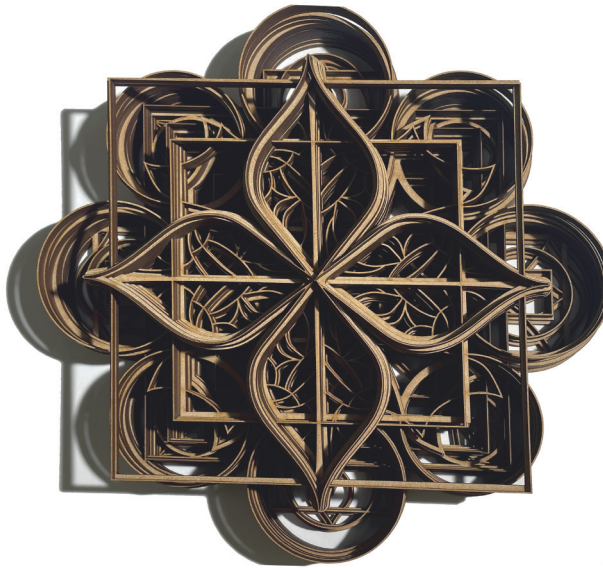




(Left) Layering by hand. (Right) Print model out as individual layers.



Order and glue layers to form a solid model.





Laser cut model from Midjourney image.

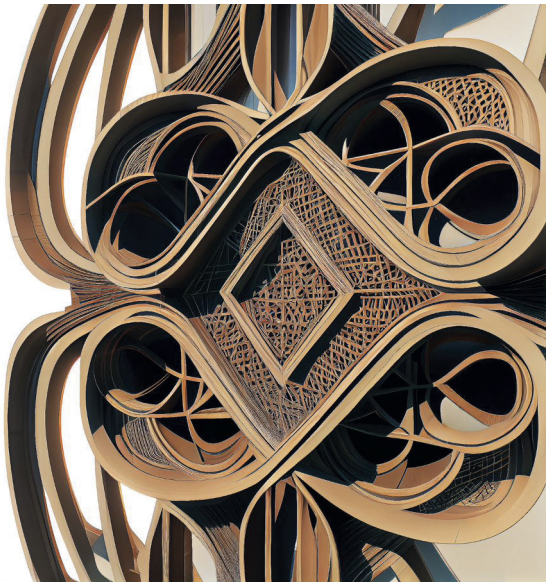
STEP 3

Laser model to Midjourney
Finally, insert a photograph of the model back into Midjourney as a reference. This process allows you to start imagining the design as a habitable space and considering how it might be built.

For the image on the next page, use the result of the image on this page as a prompt, and add the following text: 'a beautiful curving, heavily inhabited building filled with people, showing interior and exterior spaces, with lots of people inhabiting the curves on different floors and levels, lots of varied lighting, sky scraper, masterplan building --ar 9:16 –iw 1.5'.

This process could have been accomplished with a complex Grasshopper script to convert raster to vector and extract efficient linework, which might be considered more innovative. However, this process embraces the human touch amidst a generative technological craze. There is a sense of freedom in returning to analogue working methods and allowing the AI to assist in ideation.

The AI software provides continuous inspiration, while switching back to modelling helps keep the project grounded in a physical design space. This back-and-forth process expands the conceptual process by offering frequent opportunities to express design ideas. This approach seems to accelerate the design thinking, allowing more time to push concepts further. At times, unexpected results may appear, but this creates an open dialogue between your intuition as a designer and the immensely expansive feedback of artificial intelligence.



Insert a photograph of the model back into Midjourney.

Scan for full workflow GIF.



Add scale, context and depth. Midjourney v6.

TIPS

Keep an open mind	AI is constantly evolving and could one day become sentient. Be ready to adapt and evolve alongside it, and treat AI not just as a tool, but as a collaborator that will guide your creativity and innovation.
Align AI with your passions	You can train AI in whichever field you like! Find your niche and utilise AI's capabilities to enhance your unique, creative identity.
Think beyond the digital realm	AI-generated designs hold the most value when realised in the physical world. AI should augment your creativity, not replace it, so the final product reflects your vision and not just an algorithm's capabilities.

What does the future role of the architect or designer look like?

The future architect or designer is a hybrid thinker, who does not replace the human touch, but enhances it. The role will be about staying adaptable, using tools like AI to explore new possibilities, and responding to real-world needs - whether environmental, cultural, or personal. They will understand how to bring ideas to life in ways that resonate with people. In my university project, I used AI to generate forms and translated them into physical 3D models, showing how the digital and physical worlds can work together. This is the future: architects who embrace innovative technologies, while fulfilling their responsibilities to the environment and society.



Pixelstorms

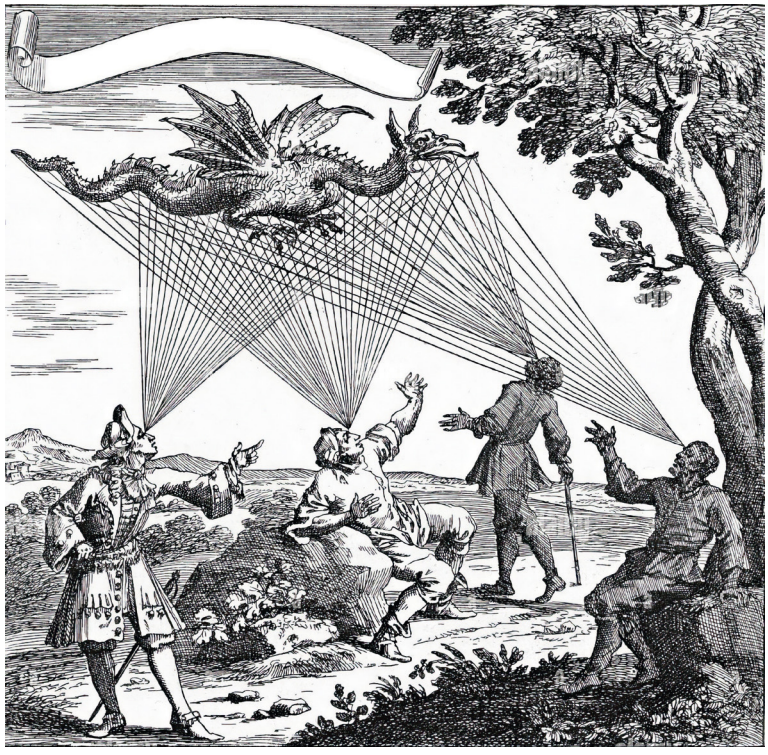
Matias del Campo

DOI: 10.4324/9781003726500-4

Ways of Seeing

The classical Greeks had some peculiar thoughts on the inner workings of human perception. There were two rival schools of thought. While the atomists, under the leadership of Democritus, proposed the theory of images (εἴδωλα) streaming from objects, the extramissionists believed that our eyes emanate rays of light, scanning the objects we perceive, thus making them visible.[1][2] A theory held close to heart by thinkers such as Plato, Euclid and Ptolemy.[3]

Imagine: total darkness, a world without light and colour, filled with objects and matter, only waiting for the human to show up to cast their gaze – formed of light rays – emanating onto the object to reveal its presence. This is known as the emission theory and there is hardly a concept more anthropocentric than this, where without humans there is nothing and the world is in eternal darkness. As usual, the earlier, older atomists, such as Democritus, were right all along with their intromission theory. As with a world made of atoms, they are right that we perceive the image of an object because it bounces light rays into our eyes as opposed to emitting from it, which we understand to be correct today. Perception based on the emission theory is not entirely alien. It draws similarities with the echolocation of bats or the sonar of whales. This system of perceiving the world is based on acoustic rays bouncing off objects in a distance; it is, however, a rather dull, grey world void of light, shade, tone and colour.[4]

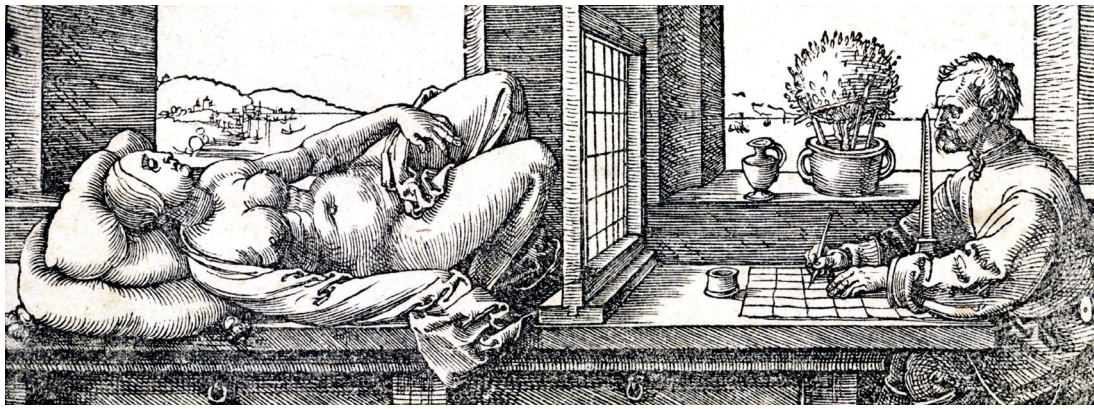


Plato, Euclid and Ptolemy believed in the extramission concept of perception.



The mimesis of existing objects. While projecting estranged properties such as the flowing veils, the underlying concrete structure of this lakeside building can be detected as a defamiliarised form of modern architecture.

Why am I interested in ancient concepts of perception in an essay about the mechanisms behind AI image generators and 3D objects, you ask. Just think of the good old ray tracer. The relationship between concepts of perception and concepts of representation of what we perceive is crucial to understand.

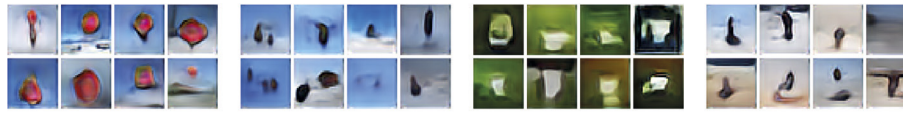


Albrecht Dürer's illustration of ray tracing.

Ray tracing, a term familiar to anyone who has toiled away using a computational rendering engine of decades past, has its origins in the concept of projecting 3D objects onto flat surfaces. It was first described by Albrecht Dürer in 1526 who is credited for its invention.[5] Among the various techniques described by Dürer are two particularly interesting ones: the concept of projecting geometry onto an image plane, as utilised in rasterisation today, and a second one that determines which kind of geometry is visible along a given ray, 'ray tracing': a concept that naturally lends itself to a translation into the mathematical ecology of computer-generated images. This was first accomplished by Arthur Appel in 1968, who devised a method that traces a ray through each point of a grid and, in doing so, calculates the closest surface to the camera, rendering (or rather 'shading') the surface visible.[6] A non-recursive ray tracing algorithm called 'ray casting' facilitated tracing secondary rays to the light source and thus helped determine which points were in the shadow and which ones were not – voilà: shadows! Many improvements followed: Goldstein and Nagel's '3D Visual Simulation', Scott Roth's flip book animation, Turner Whitted's anti-aliasing concept and so on.[7][8][9]

Jump ahead a couple of decades and a completely novel paradigm emerges from the primordial soup of computer-generated images: diffusion models. Of course, there was an evolution to get to the concept of diffusion models, but properly illuminating this episode would fill an entire book. The origins of diffusion models as we know them today can be found in research pertaining to automated image captioning.

Around 2015, researchers found that there were enough annotated images in their dataset to consider the possibility of creating an algorithm that could describe the content of an entire scene instead of just performing object recognition.[10] At that time the dataset known as ImageNet, for example, had collected around 1.3 million images (compared to today's 14 million).[11] The following year, research emerged to allow a reversal of the process; therefore, instead of using the algorithm in an analytical role, it was used to generate the outcomes itself.[12] (Sidenote: reversing the flow of information is a common technique in the exploration of machine learning methodologies, as exemplified in 'deep dreaming'. Instead of having the algorithm describing the scene, the author instead types in a text describing the scene, and the algorithm generates the image.)



First attempts at text-to-image generation as depicted in 'Generating Images from Captions with Attention'.

This led directly to further significant developments of diffusion models in the same year. In particular, they addressed a core problem in machine learning: modelling complex datasets utilising exceedingly flexible families of probability distributions that allow the sampling, learning, evaluation and inference for analytically and computationally robust results.[13][14] The concept of diffusion is based on non-equilibrium statistical physics in that it slowly and systematically destroys the structure in a data distribution by increasingly adding a forward diffusion process.[15] Therefore, reversing the process restores the structure of the data, subsequently generating a highly flexible and tractable learned path of the data. One of the defining characteristics of this approach is the speed of learning, sampling and evaluation of probabilities within the deep generative model that contains thousands of layers (time steps). In a concurrent fashion, the model can in addition compute posterior and conditional probabilities.

Thank you for bearing with me as we hike through this potentially boring valley of history, visualisation concepts and image generation technologies. Now for the more exciting part. So, what do perception, ray tracing, meaning, language, pixels and diffusion models have in common? They share the problem of representation.



Stage design for the play *'Penthesilea'* by Heinrich von Kleist, the result of prompting a diffusion model.

Representational Boundaries

Representation is a peculiar and somewhat open-ended term as it can mean so many different things depending on which discipline is seeking to represent phenomena. It plays an important role inside Western traditions that are concerned with the relationship between form and reality, as for example encapsulated in the concept of 'mimesis'. In short, its philosophical structure is made up of three elements: the world is present, humans perceive it through their sensorial apparatus and finally it is represented again in the form of a medium. As illustrated at the beginning with the 'extramissionist' idea, there are concepts of perception that include the question whether the world exists without human perception. However, I would rather dismiss this idea as a profoundly anthropocentric concept that appears increasingly unlikely the more we know about the nature of perception and intelligence.

Dogs, cats, canaries, bees and cockroaches possess perception too – do they (like humans) have individual perceptions of the world? It is more likely that neuroscientist Anil Seth is right when he poses that each individual human hallucinates their own version of the world as perhaps other various intelligent ecologies do. It could be that the small section where these hallucinations intersect is what we call 'reality'.

However, representation does not necessarily conclude in a depiction of reality. As Michael Young pointed out in his book *Reality Modeled After Images*: 'Some positions see this process as a sequence that moves from reality toward representation, each step further away from contact with the "real" as an ideal truth. In this, images are twice removed, and thus not to be trusted.'[16] Let me give you two possible philosophical interpretations that deal with this problem. One that addresses the problem of representation and rendering and is anchored in the poststructuralist world of Jacques Derrida, and another one that deals with the efficacy of language as fiction in a theory by Peter Eisenman. Probably the only 'philosopher – architect' collaboration that resulted in a built project: the garden at the Parc de La Villete.[17] Both of them are firmly in the deconstructivist camp, and considering the destructuring of data in diffusion models, we might be on to something here.



Vincent van Gogh, *Shoes*, oil on canvas, 37.5cm × 45cm, 1886.

Derrida in his book *The Truth in Painting* picks up the gauntlet that Martin Heidegger threw regarding the aspects of representation in one of van Gogh's paintings.[18] [19] I found different names for the painting, but for the sake of simplicity let us just call the painting *Shoes*. (Sidenote – Vincent van Gogh is a specifically fitting choice of an artist for an article about AI, as he is a very popular victim of image-based AI experiments – see, for example, techniques such as style transfer or Deep Dreaming). Derrida states: 'Yes, hallucinations and ghosts. But what does that mean? There are many reasons to distrust the concept of "projection" as manipulated by Shapiro… Does this concept of projection belong to objectivism, or at least to a system of subject-object relations or of truth (as adequation or as unveiling, as presence in representation or in presentation) which can no longer stand up to the effects of the unconscious, as determined by a certain state of psychoanalytic theory?'[20] Derrida is onto something here. Almost like a premonition of the ongoing debate surrounding diffusion models, which interrogate storms of pixels with de-structured data to form evocative images, shaped by our understanding of psychology and, ultimately, neuroscience. Is this process subjective? Objective? His observation of the value of hallucination



Representation of Gothic architecture through the lens of a diffusion model. A pixelstorm based on the destructuring of data capable of performing denoising along a Markov chain.

in representation – particularly in van Gogh's paintings – challenges Heidegger's notion, wherein the philosopher projects his own referential system onto the artwork, interpreting it as merely a pair of peasant shoes. This critique opens a fertile canvas for exploring the interplay between projection, interpretation and representation, questioning the boundaries of meaning in both art and artificial generation.

The technologies described above, ray tracing and rasterisation are all prone to swing the pendulum into one or the other direction: part projection (of rays or cultural values); part representation (pixels on a screen, political participation).

Imitation, Simulation and Meaning

I promised a second philosophical position regarding aspects representation, rendering and reality. All of which are worthwhile aspects to discuss in a world where technologies capable of imitating the cognitive abilities of humans are at play (i.e. memory, attention, perception, language, reasoning, metacognition etc.). From image generators based on diffusion models to large language models (LLMs) operating with pre-trained transformers, the results are the same: imitation of aspects rooted in human cognition. Machines capable of imitating the general dynamically integrative function of performing selective responses to meaningful changes in the external environment present in human cognition.[21] Rendering a representation of reality, of course, brings us directly to aspects of fiction. Eisenman discusses this problem in the light of representation as fiction that simulates (or imitates – as in the case of AI models) meaning. He poses that the first form of fiction is 'representation'. He goes on to claim that the pre-Renaissance era was defined by a consensus of representation and language. A direct translation of language to meaning with little space for interpretation. In other words, the way language produced meaning 'could be represented within language'.[22] In terms of architecture, the consequence of such a language culture is that of a Gothic or Romanesque cathedral. Historic architecture embodies the fact that truth and meaning were self-evident.[23] This shifted in the Renaissance era. To this very day buildings are defined by the simulacra of existing examples, giving value to the architecture by representing previously treasured specimens. A closer look at the buildings you love may reveal this. To this extent, the simulation of architecture represents a form of fiction. This brings us directly back to the various generative AI models that we see in action today.

This essay is an attempt to unfold, admittedly in an incomplete and crude way, the ways that perception, representation methods (technical and intellectual), imitation and meaning form a system of concepts that are connected by gravitational forces pulling them into strange irregular orbits. Connecting insights from the neuroscience of perception, technological development and the large historical trajectories of the architecture discipline, I attempt to unfold its meaning within the context of artificial intelligence and architectural image production. In effect, providing a platform to interrogate the intellectual merit of utilising diffusion-based generative AI models and large language models using transformers as a base for the interrogation of the architectural project.

Notes

1. K. Rudolph, 'Democritus' Perspectival Theory of Vision', *Journal of Hellenic Studies*, 131 (2011), pp. 67–83.
2. A. Raftopoulos, 'Lessons from the History of the Concept of the Ray for Teaching Geometrical Optics', *Science & Education*, (2011).
3. J. E. Cottrell, & G. A. Winer, 'Development in the understanding of perception: The decline of extramission perception beliefs', *Developmental Psychology*, 30(2) (1994), [doi: 10.1037/0012-1649.30.2.218], pp. 218–228.
4. B. Fenton, & F. Jensen, E. Kalko, & P. Tyack, 'Sonar Signals of Bats and Toothed Whales', *Springer Handbook of Auditory Research*, 5 (2014), [doi: 10.1007/978-1-4614-9146-0_2], pp.11–59.
5. P. Parshall, 'Graphic Knowledge: Albrecht Dürer and the Imagination', *The Art Bulletin*, 95(3) (2013), pp.393–410. JSTOR, http://www.jstor.org/stable/43188839, accessed 27 Feb 2025.
6. A. Appel, 'Some techniques for shading machine renderings of solids', *AFIPS '68*, Proceedings of the spring joint computer conference, 30 April to 2 May 1968, [doi: 10.1145/1468075.146808], pp. 37–45.
7. R. Goldstein, A'Nagel, '3-D Visual simulation', *Simulation*, 16(1) (1971), [doi: 10.1177/003754977101600104], pp. 25–31.
8. S. Papert, *Mindstorms: Children, Computers, and Powerful Ideas*, (New York City, NY, 2001).
9. T. Whitted, 'An improved illumination model for shaded display', *SIGGRAPH '79*, Proceedings of the 6th annual conference on Computer graphics and interactive techniques, [doi: 10.1145/800249.807419], p. 14.
10. O. Vinyals, A. Toshev, S. Bengio & D. Erhan, 'Show and tell: A neural image caption generator', 2015 IEEE *Conference on Computer Vision and Pattern Recognition* (*CVPR*), (2014), [doi: 10.1109/CVPR.2015.7298935], pp. 3156–3164.
11. J. Deng, W. Dong, R. Socher, L-J. Li, K. Li & F. F. Li, 'ImageNet: A large-scale hierarchical image database', 2009 IEEE *Conference on Computer Vision and Pattern Recognition*, (2009), [doi: 10.1109/CVPR.2009.5206848], pp. 248–255.
12. E. Mansimov, E. Parisotto, J. Ba & R. Salakhutdinov, 'Generating Images from Captions with Attention', *ICLR*, (2016), [doi: 10.48550/arXiv.1511.02793].
13. J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan & S. Ganguli, 'Deep Unsupervised Learning using Nonequilibrium Thermodynamics', International Conference on Machine Learning, (2015), [doi: 10.48550.arXiv.1503.03585].
14. C. Ley, S. Babić, & D. Craens, 'Flexible Models for Complex Data with Applications', *Annual Review of Statistics and Its Application,* 8 (2021), [doi: 10.1146/annurev-statistics-040720- 025210].
15. Y. Bahri, J. Kadmon, J. Pennington, S. S. Schoenholz, J. Sohl-Dickstein S. Ganguli, 'Statistical Mechanics of Deep Learning', *Annual Review of Condensed Matter Physics*, 11:1 (2020), pp. 501–528.
16. M. Young, *Reality Modeled After Images: Architecture and Aesthetics after the Digital Image*, 1st edn., (Oxford, 2021), p. 4.
17. B. Wocke, 'Derrida at Villette: (An)aesthetic of Space', *University of Toronto Quarterly* 83 (2014), [doi: 10.3138/utq.83.3.739], pp. 739–755.
18. J. Derrida, *The Truth in Painting*, tr. G. Bennington & I. McLeod, (Chicago, 1987).
19. I. Thomson, 'Heidegger's Aesthetics', The *Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta & U. Nodelman, 2nd edn. (2024), https://plato.stanford.edu/entries/heidegger-aesthetics/, accessed 3 March 2025.
20. Derrida, *The Truth in Painting*, pp. 366–367
21. A. Abraham, *The Creative Brain: Myths and Truths* (Boston, 2024), p. 146.
22. Peter Eisenman in K. Nesbitt, *Theorizing a New Agenda for Architecture: An Anthology of Architectural Theory, 1965-1995*, 1st edn., (New York NY, 1996), p. 213.



Real-Time Façade Iterations

ComfyUI + Grasshopper

Ömer Nuray

DOI: 10.4324/9781003726500-5

JOURNEY

Real-time visualisation with AI emerged from a year of intensive research aimed at answering a central question: How can we minimise the friction between AI and architectural conceptual design? Since my first AI-generated image, I've been dedicated to answering that question by making AI both controllable and genuinely useful for designers. Among the various AI tools, Stable Diffusion stood out as the most promising one for a seamless integration. Its open-source community and wide range of compatible extensions clearly set it apart. ControlNet, its most notable feature, allows designers to guide image generation using diverse inputs such as sketches and 3D views, making it an invaluable tool for accurate design.

My initial step towards real-time design was linking one of my Grasshopper scripts to Stable Diffusion, allowing for a direct visualisation of my parametric model. While this reduced friction, image generation remained slow and lacked a true 'real-time' feel. However, recent advancements in generative AI models have dramatically changed that. Models like SDXL-Turbo and SDXL-Lighting provide a smooth, interactive, real-time experience, instantly reflecting changes made t o the 3D model, bringing us closer to true real-time visualisation.

Generative AI models excel in broad generation capabilities and remarkable creativity, which is both their greatest strength and their most significant challenge. By integrating these models into our individual design processes, we can harness AI systems as powerful tools that complement our thinking rather than compete with it.

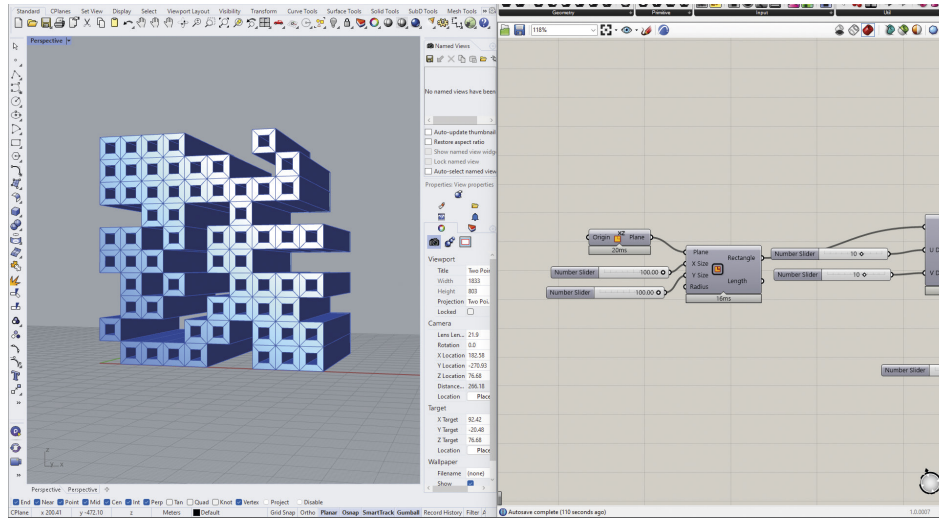
Once a design direction is established through this real-time visualisation process, I also use AI to generate a full-resolution, highly detailed image of the final concept. Below, I'll outline step-by-step how to implement this workflow into your own design process.

WORKFLOW

STEP 1

Preparing the 3D model

The first step is deciding which 3D model will serve as the input for our real-time workflow. This workflow's versatility opens doors to virtually any kind of 3D model you can imagine, including physical maquettes captured via webcam. For optimal results, however, it is recommended to use a parametric model created in Grasshopper. This approach unlocks the full potential of the integration, as parametric models offer greater flexibility and enhanced real-time manipulation capabilities.

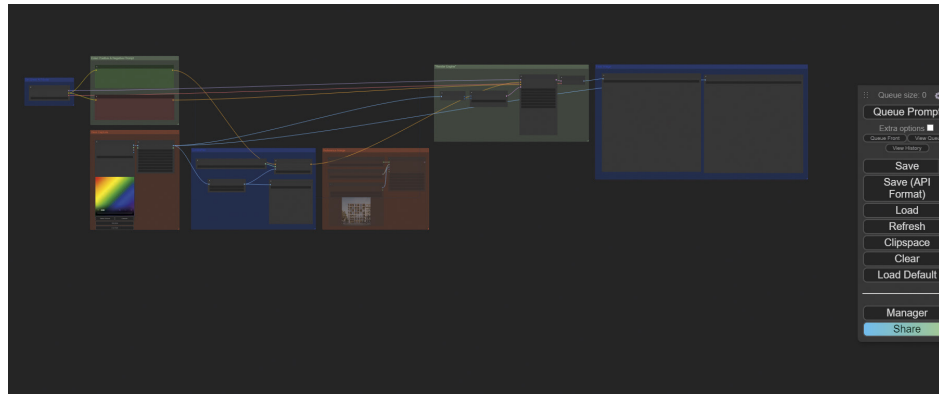


Preparing the 3D model.

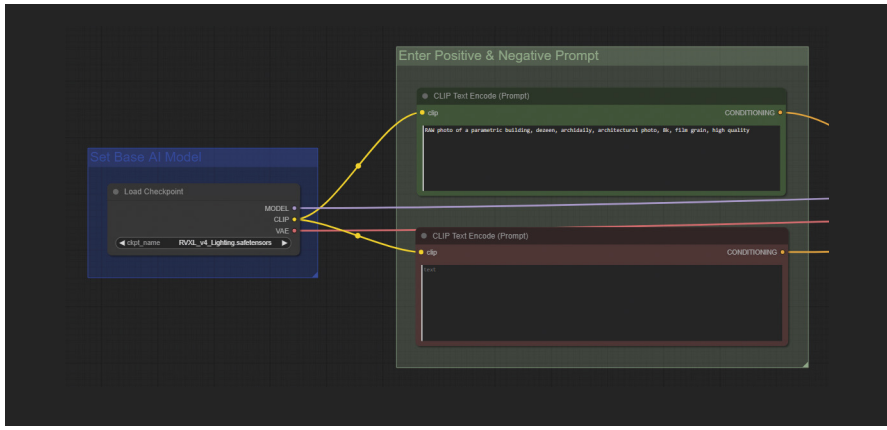
STEP 2

Building the workflow in ComfyUI

Next, we'll create the workflow inside ComfyUI for Stable Diffusion. We'll build upon the core 'text-to-image' process by layering additional inputs, aiming for the best results in terms of quality and accuracy. For visual clarity, we can divide the workflow into colour-coded groups based on the specific functions they perform within the process.



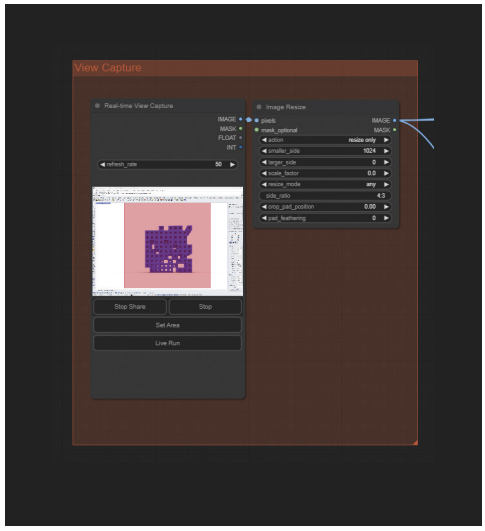
Creating the custom real-time render workflow in ComfyUI.



Setting checkpoints and defining positive/negative prompts.

Start by loading your desired checkpoint. The checkpoint is the main AI model where all the 'learned' information is stored to be used in new generations – it's one of the crucial elements in the process. The AI model you choose will have a significant effect on the final results.

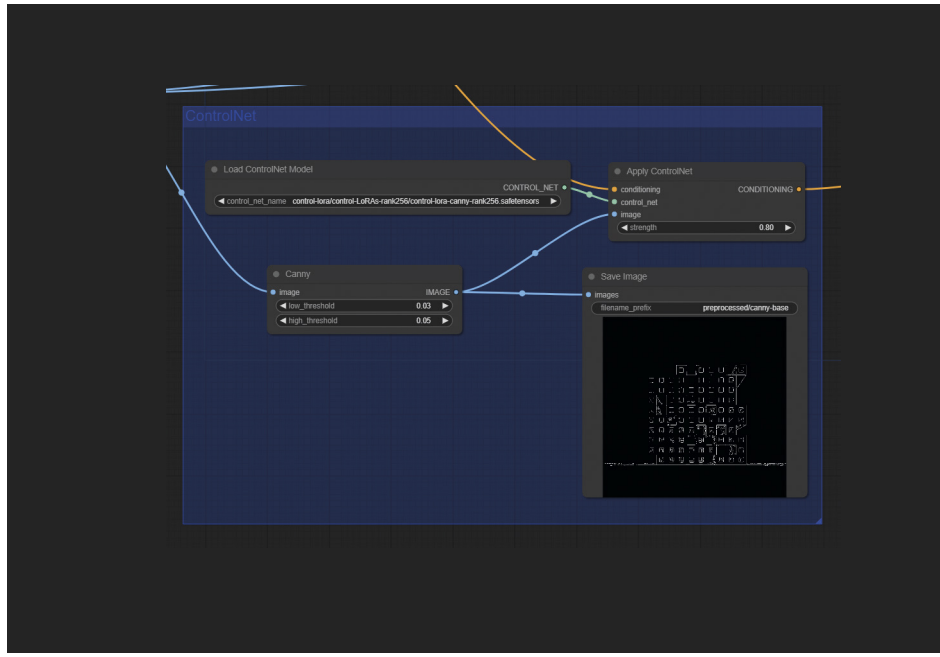
Next, describe the scene for the new view you envision in the 'Positive Prompt' text box. It's better to start with a simple prompt and add details as you refine the image. Focus on key elements like the surrounding environment, architectural style, materials, time of day, lighting and even inspirational designers or architects. There's no magic formula for prompting; it will vary greatly depending on the checkpoint used. What's more important is understanding the logic behind it – you're communicating your ideas to a machine, not to a person, colleague or architect. With practice, you'll soon understand how different prompts affect the design.



Capturing synced views between 3D modeling software and Stable Diffusion.

Initial prompt: ‘RAW photo of a parametric building, dezeen, archdaily, architectural photo, 8k, film grain, high quality.’

Use the ‘Screen Share’ component to synchronise your 3D model view (Rhino 3D) with ComfyUI. Enabling Stable Diffusion to receive a live feed will allow you to see changes in your model instantly reflected in the generated images.



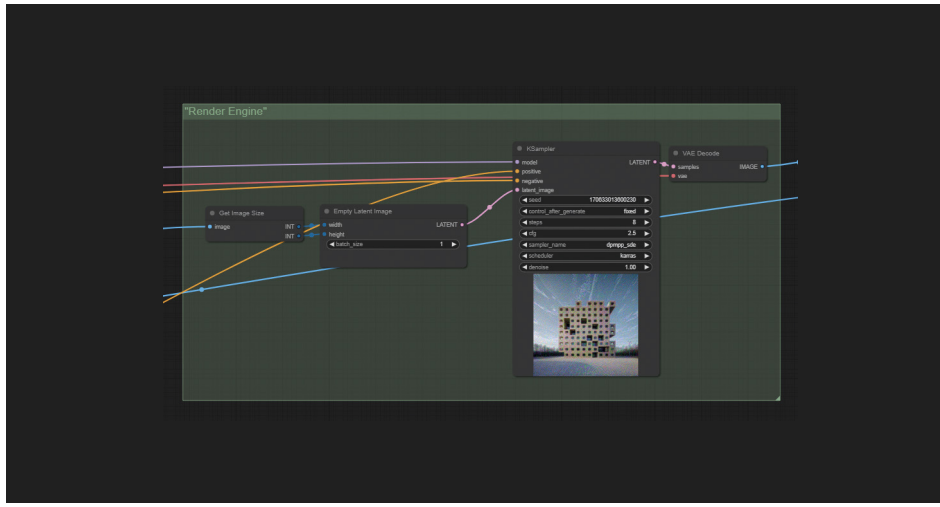
Connecting the 3D model view to ControlNet.

As stated before, ControlNet is the essential element within this process. It works with Stable Diffusion to guide image generation using existing shapes and forms. Simply put, it helps refine and develop the final image based on the initial image geometry.

In the example shown, we use the Canny model within ControlNet to provide an edge map – essentially an outline or stencil that traces over the provided image. This layer of control helps achieve accurate results that closely reflect the actual design, ensuring the AI-generated visuals maintain the core structure of your 3D model.

The final step is to connect all the inputs: the text prompt, checkpoint, the Canny edge map (generated by ControlNet) and an empty latent image. These elements are then fed into the KSampler.

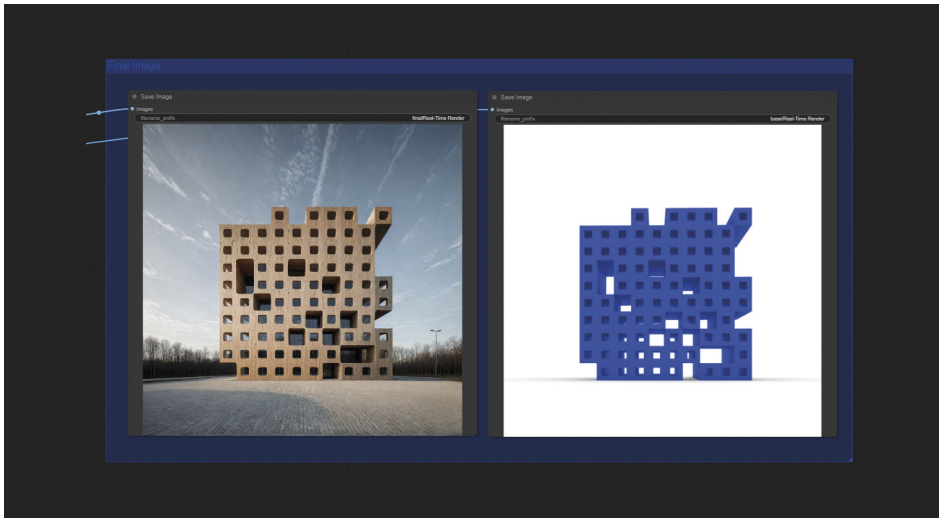
The KSampler is the heart of the image generation process. It’s where the design takes shape, with the image being generated step-by-step using the diffusion method. This process can be controlled with parameters such as seed, number of steps, classifier free guidance scale (CFG), denoise and sampler/scheduler types.



Linking inputs to the KSampler node for the diffusion process.

For this generation, the following values were used: Seed: 170633013600230, Steps: 8, CFG: 2.5, Sampler: DPMPP_sde, Scheduler: Karras, Denoise: 1.

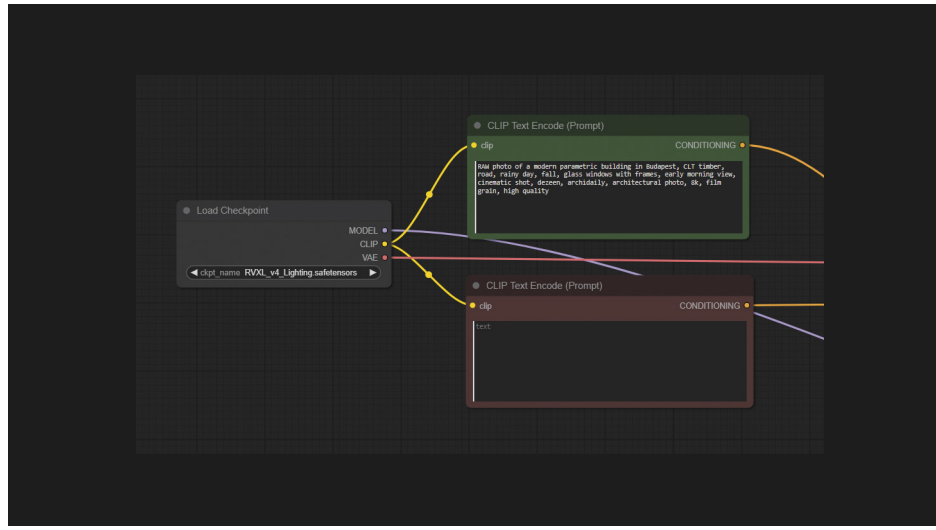
At the final step of the workflow, we can see the base 3D model view and the final AI-generated render side-by-side for comparison, allowing us to evaluate how accurately the AI has interpreted and enhanced our design.



Before and after, initial results.

STEP 3

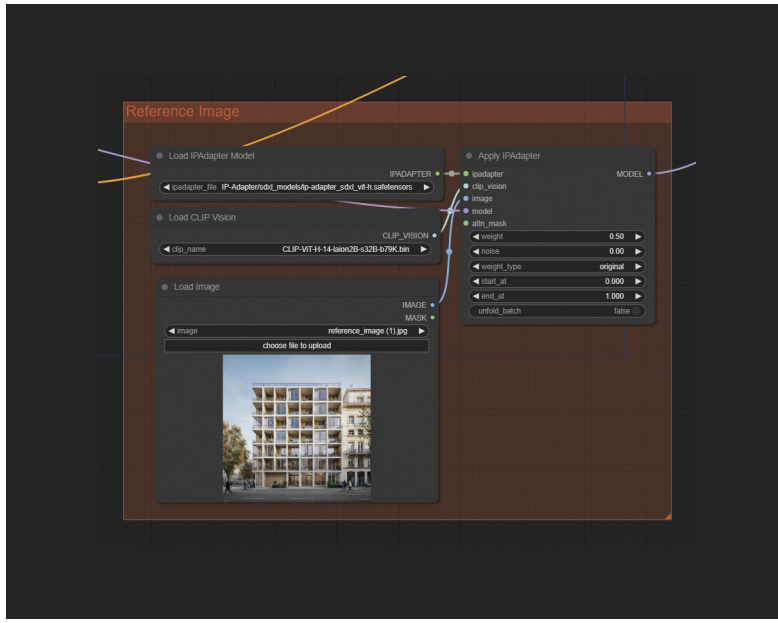
Refining the AI prompt
Once the workflow is completed, begin refining the prompt. Add details about the important aspects of the design, focusing specifically on elements missing from the image compared to the initial prompt. If a material type described in the prompt isn't visible in the generated image, try describing it in different ways or include more specific details. This iterative process helps fine-tune your design vision.



Refining the prompt and adjusting settings.

STEP 4

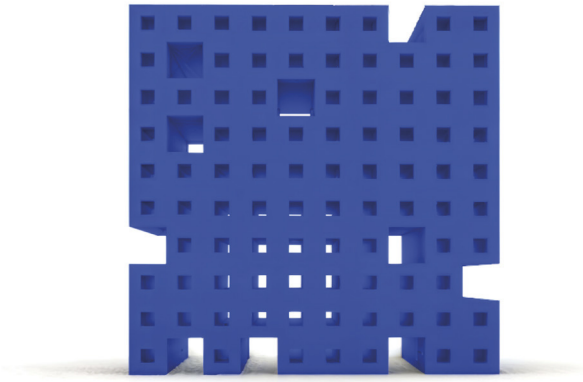
Adding a reference image
Reference images can be used to influence image generation. They are an effective way to achieve a specific material or style, or replicate the ambiance of an image. Using reference images is often more impactful than a thousand words in prompts, so leverage them to your advantage and save time on prompting. In addition to improving accuracy with certain materials or styles, using reference images often enhances the overall quality of the final image.



Adding a reference image.

STEP 5

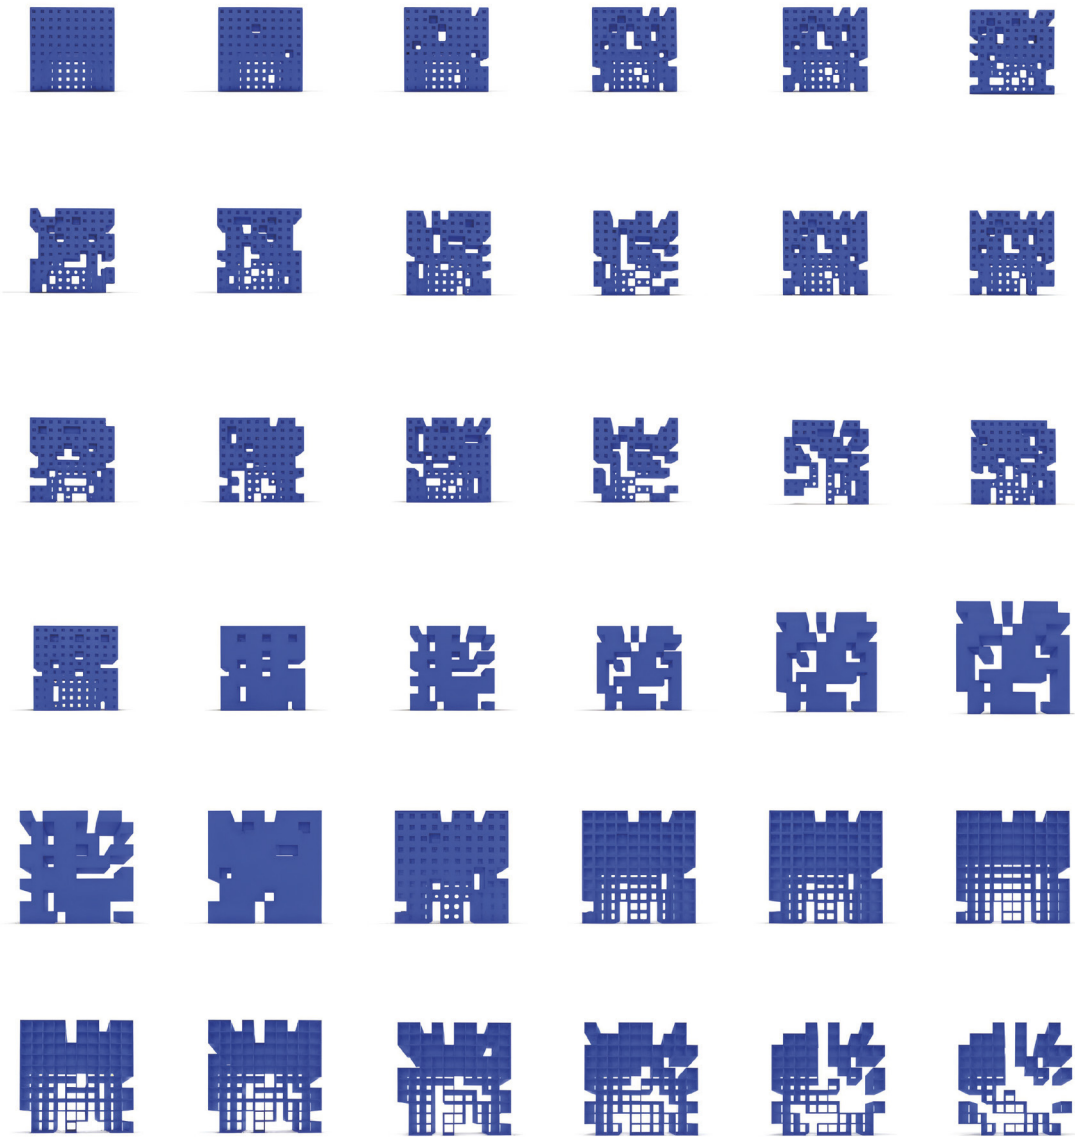
Editing the 3D model for real-time visualisation -
With the addition of the reference image, this workflow is now complete. Activate ‘Live Run’ for the View Capture node and ‘Auto Queue’ under ‘Extra’ options in the main ComfyUI menu to feed real-time synced views from the 3D model. As you make changes and navigate within the 3D model, you’ll see instant updates reflected in a realistic render view. -



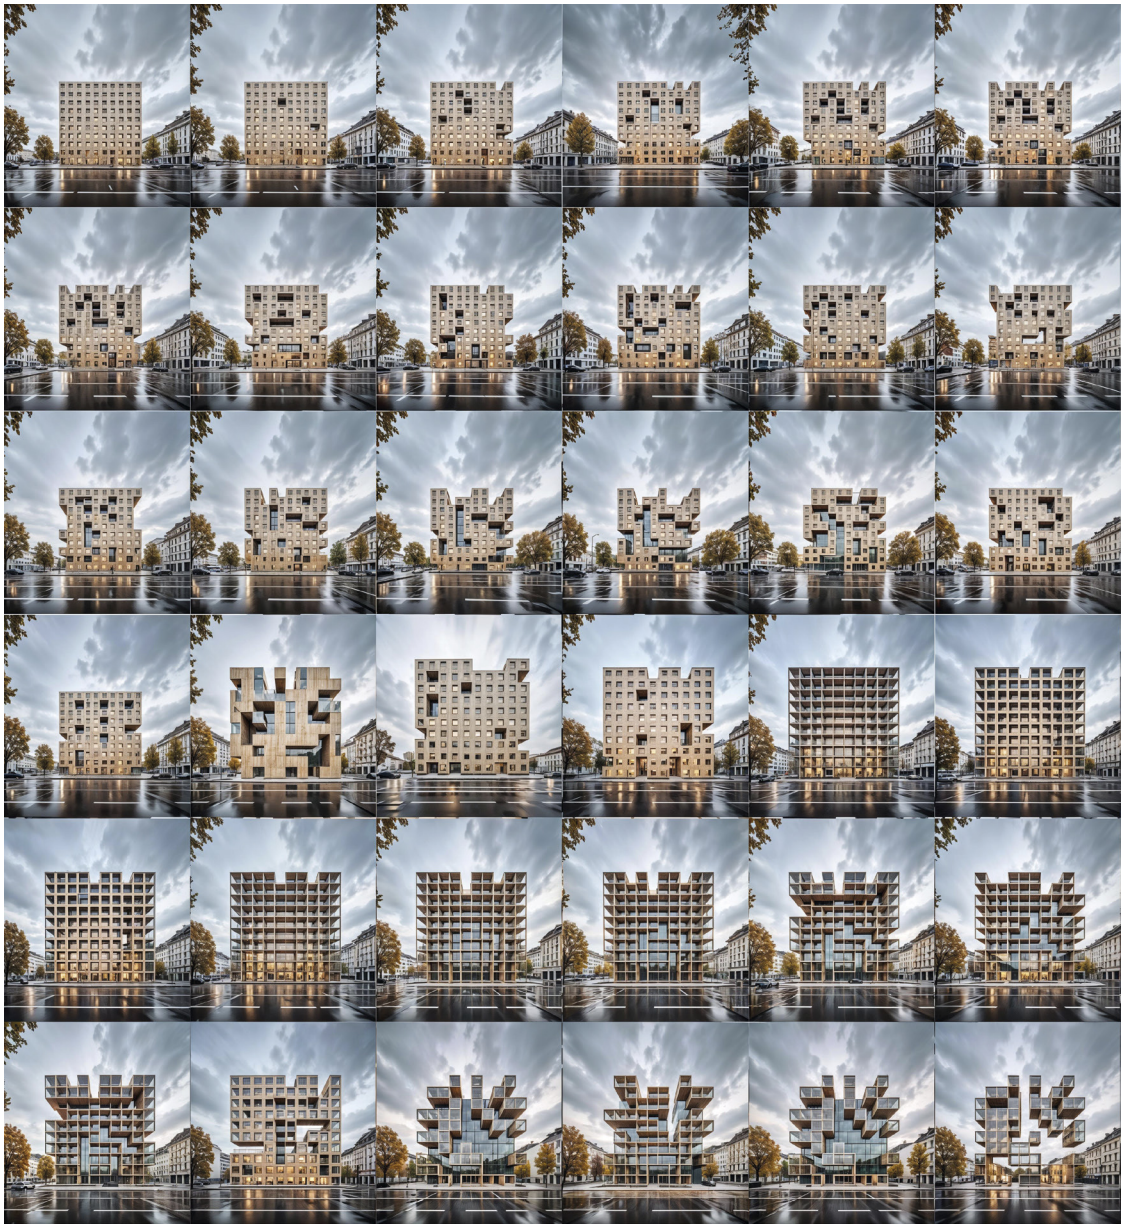
Editing the 3D model
based on the design intention.



Visualising edits in real time.



Multiple edits to the 3D model.



Consistent real-time rendering.

The final step of the process involves making edits in Photoshop to clean up small imperfections in the image and adjust the colour tones.

Positive prompt: ‘RAW photo of a modern parametric building in Budapest, CLT timber, solid facade, road, rainy day, fall, glass windows with frames, early morning view, cinematic shot, dezeen, archdaily, architectural photo, 8k, film grain, high-quality.’

Steps: 45, Sampler: DPM++ SDE, CFG scale: 12, Seed: 170633013600254, Size: 1024x1024, Model: JuggernautXL_v9, Denoising strength: 1.0, SD upscale: 4x_NMKD-Siax_200k.pth, Upscaled Image Size: 4096x4096px.



Final AI-generated image before upscaling.



Post-upscaling - final image with Photoshop and RAW enhancements.

TIPS

Treat AI as a tool, not a replacement	Just like Rhino or Archicad, AI systems are instruments to enhance our creative process. This mindset prevents unrealistic expectations and leads to more effective integration in architectural practice.
Start with clear design intent	Never expect AI to generate ideas from scratch. Instead, approach it with well-defined goals and use it to explore variations or visualise concepts you've already developed. I've found this especially crucial when using tools like Stable Diffusion for design visualisation.
Focus on understanding system logic rather than collecting 'perfect prompts'	Learning how AI processes information and responds to different inputs is far more valuable than memorising keywords. This knowledge helps achieve consistent results across various projects.

What does the future role of the architect or designer look like?

I envision the future role of architects/ designers to be centered on decision-making rather than manual execution. Currently, architects spend the majority of their time on repetitive, manual tasks across various project stages. As AI and automation increasingly handle these routine aspects, architects will be freed to focus on what truly matters: making informed decisions that shape our built environment. The architect of tomorrow will be less focused on technical operations and more on guiding projects through thoughtful choices about design, sustainability and human experience - leveraging AI tools while remaining the essential decision-maker who ensures meaningful outcomes.



The Bias Dilemma

Hassan Ragab

DOI: 10.4324/9781003726500-6

The Interconnectedness of Everything

I find myself engaged in a variety of activities, simultaneously doing everything and nothing. My 20-year journey with 2D and 3D software has carved out an obscure, multi hyphenated career. However, the introduction of AI into my toolkit over the past two years has led me to contemplate the interconnectedness of everything. These tools are not only powerful and enabling but also prompt a reevaluation of our approaches to learning, performing and conceptualising our work.



First image generated in Midjourney v2 – a piece of furniture made from a sofa and a china cabinet, June 2022.

While I don't believe it's important to pinpoint exactly where everything started (perhaps it was the Big Bang), something significant happened about three years ago: diffusion models became accessible to mass users. Without that development, you probably wouldn't be reading this – it's the sole reason I find myself in a position to write it.

You see, as I write this, I am not sure about what I do anymore – a sentiment I've come to see as a blessing. It feels like I can do anything. Without delving into my resume, the past two years have significantly influenced my career. While several factors contributed to this change, one stands out: how AI liberated me from technicalities.

Reflections on the Role of the Architect: A Brief History

Historically, while architectural marvels that have stood the test of time required great endeavour and a master builder to construct, most architecture (e.g. houses for people) was built by the owners themselves, with knowledge only of their environment and where to source building materials, similar to a bird building its nest. As we delved deeper into civilization, this practice became more specialised, evolving into what we now recognise as the role of a master builder, someone who knows how to build comprehensively, understanding all principles and trades necessary for constructing a building. As technology and the industry became more complex, the discipline branched into trades, architecture and engineering; all of them working together to create structures that transcend the natural capabilities of our distant non-human cousins. At this point, architects still needed to possess a broad knowledge base, though not necessarily an understanding of everything.

Further into the 20th century, the world witnessed a sea change with the advent of modernism and industrialisation. This era demanded an army of architects for construction projects, led by a master or senior architect who would guide the vision for hundreds of others. These architects now focused on creating schematic designs, drawing plans and making sketches.

With the introduction of computer-aided design (CAD), the master builder could no longer (and no longer needed to) keep up with technological advancements. The industry transformed into a collective intelligence, divided among thousands of individuals to build a project. Specialisations emerged within the field: structural engineers, site engineers, project managers, foremen, procurement engineers and architects who specialise in designing specific types of buildings. Computational designers and research and development (R&D) departments focused on advancing building technology rather than the building process itself.

For the first time in history, platforms now exist where one doesn't need to know how to build, draw, or even use computer software to visualise a building or an environment. Over recent months, many individuals from diverse backgrounds have entered the realm of 'architectural design' through image visualisations. With technical barriers disappearing, these newcomers are introducing fresh ideas and visualisations, imitating architects who use AI.

***

In 2005, many students in architecture school would have used the same tools that architects had been using for the past century for the first two years: canvas, pencils, ink pen sets, watercolours. In the following years, computer tools were gradually introduced into projects, but you still would have need to learn AutoCAD, Autodesk 3ds Max, Photoshop to name a few tools, at a time when there was no YouTube. The best chance you had was to take a private course outside of school, probably a summer course.



Hybrid Identities. Midjourney v6, March 2023

The real excitement for me didn't begin until I was introduced to a tool within Rhino 3D called Explicit History (the old name for a not-so-famous Rhino plugin, now known as Grasshopper). Grasshopper allows users to create parametric designs by connecting nodes – each representing a specific function or operation – without having to write conventional code. In other words, you can define relationships and rules (parameters) for your design, and Grasshopper will generate or modify geometry in Rhino accordingly.

It was at that moment I knew Grasshopper wouldn't only change my life but would also revolutionise the entire architectural practice. Because Grasshopper allows architects and designers to work with code visually – without necessarily having to learn traditional coding – it opens up new possibilities for parametric design. Its visual approach aligns perfectly with architecture's universal language, making the learning curve less steep. Moreover, Grasshopper provides a deeper understanding of digital geometry, enabling designers to create complex forms for both aesthetic and functional purposes. This control over geometry also propels the practice forward, particularly in automating key processes such as design management, fabrication and assembly.

From that moment forward, I knew I needed to master this tool. Despite the challenges, I had faith that this was the future, setting me on a path to become a computational designer.

Many years later in 2022, I wrote my first Midjourney prompt and my decade-long perceptions on design process drastically shifted.

From that point, I couldn't stop using AI tools. It felt as if I was standing at the entrance to a portal leading to an infinite space of ideas. I was compelled to explore more ways to develop a new 'process' – a method to unlock the potentials of this 'thing' that sometimes words can't adequately describe. It became an obsession



The City is a Tram. (Left) Midjourney, September 2022. (Right) Midjourney + Magnific AI, December 2023.

of how to achieve the best results. I developed several strategies to 'cheat' the system and overcome dataset biases, which were very apparent in these early models. This allowed me to explore different concepts and to mix everything I could think of with a freedom that I never had as an architect: from Art Nouveau styles to narratives where buildings can dance. I learned how to produce powerful visual outputs, control the mood of an image and utilise Stable Diffusion models in parallel.



Green Feathers Building.
Midjourney, September 2022.

The beauty of the old models was their imperfection; they 'hallucinated' quite often. As they struggled to predict what the user wanted (based on the prompt), they would often get lost into their predictions, making weird connections that a human being wouldn't normally create. It was this unpredictability that provided the perfect gateway to what some might call 'creativity'. I have always believed that innovation comes from rule-breaking, from challenging the status quo in some way. Stable Diffusion 1.5 and Midjourney v3 became the optimal tools for questioning, essentially, everything. It was a genuine quest for a new visual vocabulary, achievable only through uncontrolled, ambiguous exploration in a latent space. This exploration forged novel connections and relationships between concepts I had never considered before, opened by technology in a way that neither time nor traditional constraints would allow.

However, although a process was now in place and it was time to build on that foundation, it wasn't long before newer models were released (such as Midjourney v4, v5 and v6). With each model it became evident how much these models had evolved, how it's now easier to achieve a 'good' result. The models are 'hallucinating' less, becoming more like rendering software. Yet, this shift also opened new avenues for exploring novel expressions. One that I have been involved with for a while now is visual vocabulary that is based on intangible elements – a good example is collective behaviours and how they form the visual aesthetics of their paradigms. However, this also meant that I had to adjust my strategy, starting from scratch. I needed to learn how to prompt again and how to navigate new interfaces, as Collab notebooks were replaced by Automatic-1111. For someone not deeply immersed in technology, adapting to Stable Diffusion presented a challenge. Yet, the joy derived from mixing different platforms kept me exploring.



Tree House. Stable Diffusion 1.5, January 2023.

But, of course, this wouldn't last long. New breakthroughs emerged, presenting new opportunities. At this point, I realised I couldn't keep up with every new software release and needed to develop the ability to sift through the technological abyss I was about to fall into. It wasn't easy, but I managed. One tool that particularly caught my attention was WarpFusion. If you haven't heard of it, by the summer of 2023, it was the most reliable AI video tool in terms of image consistency (which, compared to the current models of January 2025, is considered ancient). WarpFusion was not easy to learn – none of them were, but this one was even harder. It was particularly challenging to create architectural concepts with it. I felt like I was the only architect out there using it and for good reason – this tool, like all the generative AI tools I had been experimenting with, wasn't designed for architectural work. Yet, the results achieved with WarpFusion were too compelling to ignore.

Then, guess what? While I was deeply engrossed in mastering WarpFusion, ComfyUI was released. It felt like the Grasshopper equivalent for generative AI image production. Naturally, this meant dedicating time to exploring its capabilities. My research using WarpFusion came to an abrupt halt with the release of AnimateDiff. It wasn't just me – almost everyone I knew working in this space shifted their focus. If you search on YouTube now, you'll likely find thousands of ComfyUI and AnimateDiff tutorials. These tools are beautifully complex and offer a range of possibilities. With the advancements in Midjourney v6, AnimateDiff and the introduction of Stable Diffusion XL, it felt like, for the tenth time in just 16 months, it was time to invest in a new process.



The Cities We Build. Stable Diffusion 1.5, March 2023.

In February 2024, while working on an animation for a client using the cutting-edge AnimateDiff, OpenAI released Sora, something nobody anticipated would arrive for at least another two years. There's much to learn and even more to observe.

***

When asked to contribute to this book, my challenge was summarising my experiences and insights with AI into a single chapter. There's much to learn and even more to observe. I understand that my role here isn't to guide you or show you how to do something specific. Instead, it's to highlight a moment in time. This moment feels akin to the introduction of CAD into practice, yet it's entirely different and exponentially faster and more impactful. In many conversations, whether at conferences, privately, or in interviews, the same question often arises: will AI replace architects?



Identity on the Verge of Collapse.
Midjourney v6, March 2023.

The question of AI replacing architects should rather focus on understanding the current role of an architect. What do we mean by 'architect'? Is it someone like Zaha Hadid, Bjarke Ingels or maybe Thomas Heatherwick? For me, that's not accurate, as 99% of architecture school graduates will probably never reach that level of recognition. In my view, the real definition of an architect is typically someone burdened with debt from their education, engaging in internships or, in the best-case scenario, working overtime and underpaid. They strive to amalgamate client requirements with data through plans, sections and schematic designs. To be in a position of influence, where your ideas and dreams could manifest, often means seeing those visions diluted and diminished by the very industry meant to nurture them.

Modern society has fostered a practice that has transformed architecture into a vast machine of space occupation with architects serving as cogs to keep this machine running. However, we now face new technology that not only produces a cog that is faster, cheaper and more capable, but also one that is numb, non-intuitive and more obedient (for now).

One doesn't need to predict the future by looking forward; observing the past and the present suffices. I have come to believe that the present is a repetition of the past – things change, and the pace of this change is exponential and time sensitive. Just as architects once relied on right angles, triangles and templates to assist in their drawings, and just as I once utilised 3D assets, templates and scripts to aid my school and professional projects, I now turn to AI cloud computation to bring my ideas to life. Perhaps one day, all I will need is to think about a design and it will come to fruition. At that point, the process – my process – will be valued not for how I implement it, but for what I have learned from it.

To tell you the truth, I don't think anybody truly knows what the future holds. In one interview, Yuval Noah Harari mentioned that for the first time in history, we as human beings are uncertain about how the future will unfold. This uncertainty makes it challenging to advise our children on what they should learn to be prepared for the future. Similarly, Jensen Huang has stated that there's no longer a need to learn traditional coding, as all future code will be in human language. Emad Mostaque, formerly CEO of Stability AI, also predicted that human coders might not exist in five years.

Here I stand, investing almost every waking hour in exploring novel ways to implement and understand these tools. I continue to develop different processes at an unprecedented rate for me (and I'm considered to be at a world-class competitive level, or at least that's what they say about me). As each process becomes obsolete, a new sort of meaning takes its place, and I feel compelled to share it with you even as I am still making sense of it myself. One thing is certain: change has always been the only constant. As I write this, the rate of change has accelerated to its peak. However, I can't predict how it might appear to you, whether you are reading this at the time of release or decades from now. Just know that while things are changing, for better or worse, this is the cycle of humanity. Both you and I are merely sailors navigating through a sea of doubts and unknowns. All we have are our processes, which may soon become obsolete.



Brutal Criminal - still frames from video animation. Vid2Vid Stable Diffusion.









‘I Prompted a Sofa, then Sold It!’

Midjourney + Rhino + Woodwork Upholstery

Rada Daleva

DOI: 10.4324/9781003726500-7

JOURNEY

My entrepreneurial journey began at the age of five, amidst the vibrant atmosphere of a seaside business fair. Inspired by my parents' stand, I set up my own, selling painted seashells. This early fascination with entrepreneurship and nature evolved into Daleva Design, a luxury brand specialising in biomorphic furniture. Our pieces, inspired by nature's organic forms, seamlessly blend art and functionality, reflecting the creativity and uniqueness that first captivated me in those seashells many years ago.

My journey from a childhood fascination with seashells to designing biomorphic furniture was not linear. It was sparked by the power of artificial intelligence. The generative AI program Midjourney became my creative partner, allowing me to explore and expand my imagination beyond what I thought possible. From the moment Midjourney was released, its capabilities captivated me. I spent countless hours experimenting with various design concepts, generating hundreds of variations daily, though only a select few were shared on social media platforms. These choices were deeply rooted by my personal taste and design preferences, helping to shape a distinctive brand. With Midjourney's help, I have been able to think outside the box and reinvent the art of integrating nature into our homes in a subtle yet impactful manner.

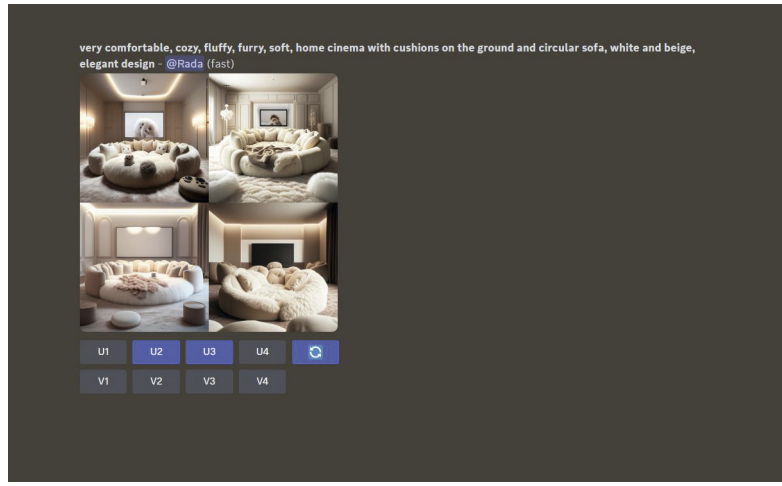
Now let me take you through the process of how to prompt a sofa in Midjourney and how it can lead to client interest, fabrication and the launch of an AI-powered design brand.

WORKFLOW

STEP 1

Prompting the sofa design

Start by articulating your vision for the sofa in simple terms, describing its intended style, functionality and aesthetic. Input this description into Midjourney, and let the AI generate a range of design concepts based on your prompt. After reviewing the generated designs, select the ones that best capture the essence of your vision.



Enter a simple prompt that describes your idea.

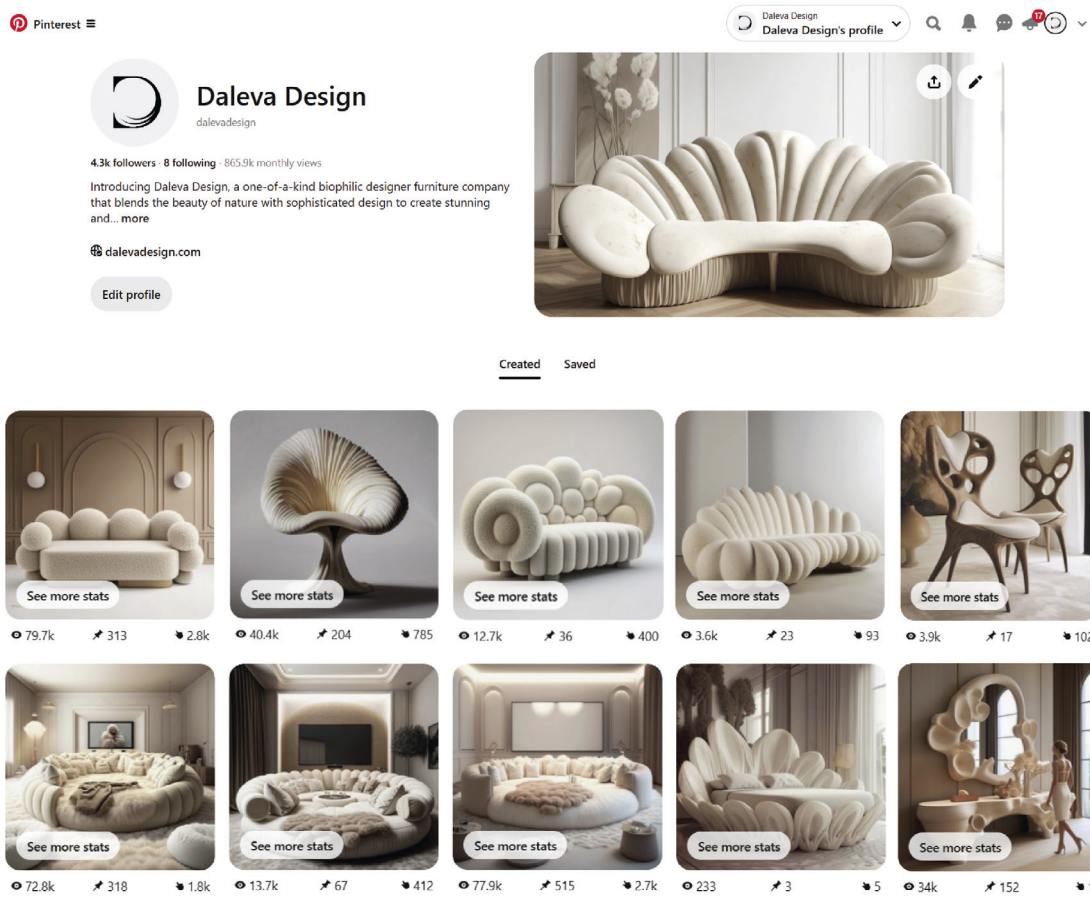


Select the best result.

STEP 2

Sharing on social media

You can share the selected designs on social media platforms, such as Pinterest and Instagram. This step will help attract public interest in the designs, gather feedback from potential clients and engage with individuals or businesses interested in purchasing the products. Ensure the posts align with and articulate your style in a professional manner that appeals to clients rather than the design community. This approach will help in establishing a recognisable brand identity and attracting a target audience interested in high-quality and, for example, custom-designed biomorphic furniture.



Sharing the designs on social media.

STEP 3

Communicating with the client
Even without a large social media following, it’s possible to generate client interest and purchasing potential for your product by adopting the right search terms. Initial contact may lead to calls and meetings where you can discuss their specific preferences, requirements and any desired modifications. Conversations are crucial for tailoring your design to the client’s needs and ensuring their satisfaction with the final product. Through this collaborative process, you may be able to reach a deal and begin the production.

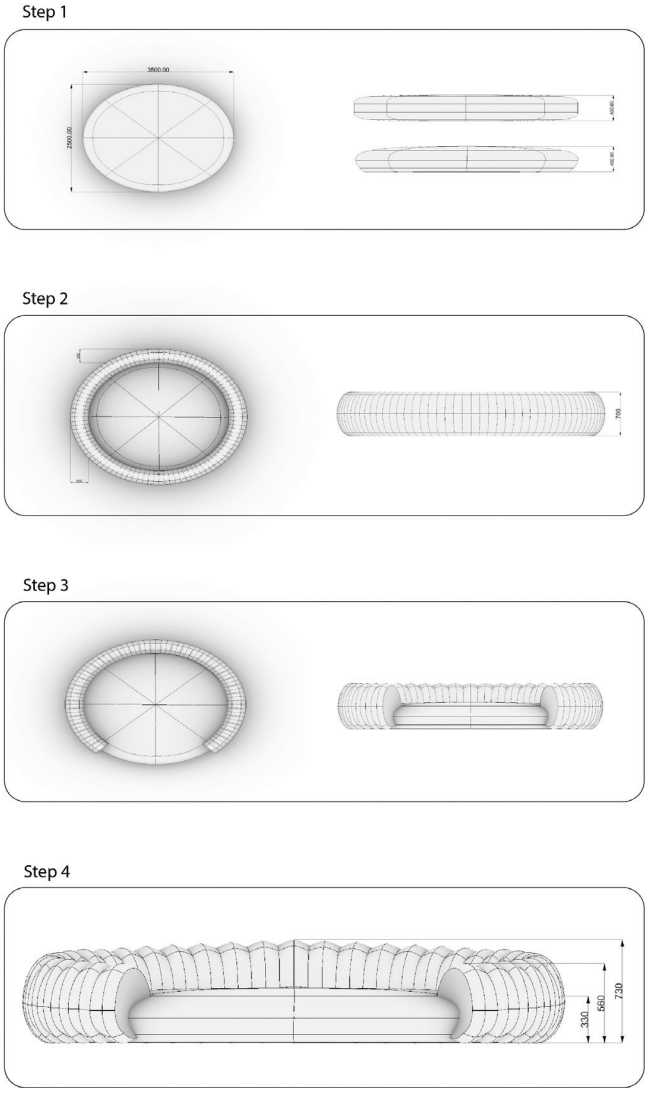
STEP 4

Creating a 3D model
Using your architectural background and software skills, translate the selected design into a detailed 3D model using Rhino. This model will allow you to visualise the sofa from various angles and perspectives, and it will serve as a tool for communicating the design to the client. Incorporate the client’s feedback into the model, adjusting elements such as size, shape and materials to align with their preferences. Additionally, the 3D model plays a crucial role in the manufacturing process by providing precise measurements and specifications, ensuring that the manufacturers can produce the sofa accurately and according to the design. Creating an accurate 3D model of a sofa in Rhino is essential for ensuring a smooth production process and delivering a final product that aligns with client expectations. By following a systematic approach, you can model a precise and visually compelling design.

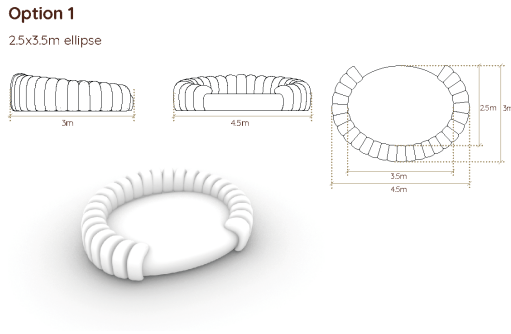
To begin, create a SubD Ellipse in Rhino that matches the sofa’s desired dimensions. Adjust the number of subdivisions in the ellipse to balance smoothness and flexibility for shaping. This will form the base of the sofa seat. Ensuring the dimensions are correct at this stage is crucial, as it sets the foundation for the rest of the model. Next, model the arms of the sofa using a SubD Torus. Define the torus with subdivisions that divide each arm into three sections: two for the sides and one to form the arch. Trim the torus to create an opening for the seat. With the torus prepared, the next step is to add definition to the arms. Select every other edge along the torus and use the ‘Crease’ command to give the arms a ridged appearance. Scale down the creased edges slightly, shaping the arms to taper inward as they extend towards the middle. To achieve the characteristic upward curve seen in the AI image, focus on the middle subdivision of each arm and its adjacent section. Gradually increase their height, ensuring the adjustments are mirrored on both sides for symmetry. Continue refining this curve until the arms reach their highest point in the middle, creating a sleek and elegant design. After shaping the arms, verify all dimensions to ensure they align with the design specifications. Double check the proportions, symmetry and

alignment of all components. If any adjustments are necessary, fine-tune the model by manipulating control points or subdivision edges. Finally, prepare the 3D model for manufacturing. Add any necessary annotations or measurements to the file for clarity, ensuring that manufacturers can easily interpret the design. Export the model in a format compatible with production requirements, such as .stl, .obj, or .step.

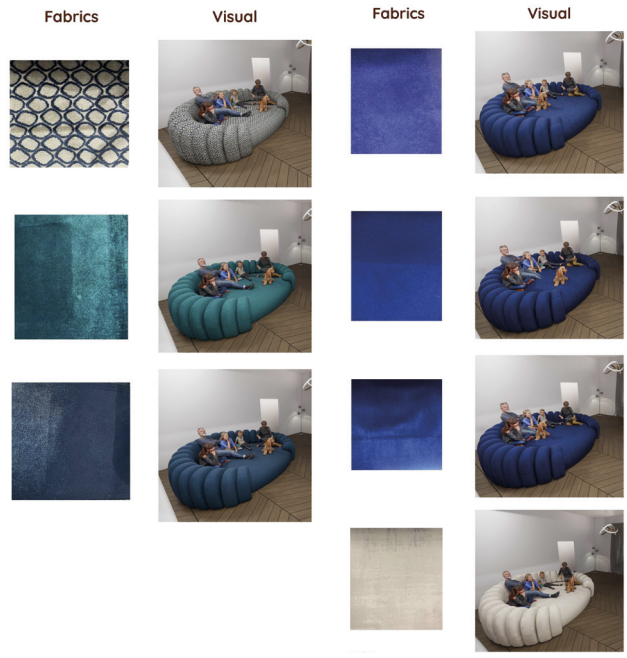
By following these steps, you can create a precise and detailed 3D model of a sofa designed with AI that translates seamlessly from design to production, ensuring a final product that meets both your expectations and those of the client.



Building a Rhino model of the sofa.



Discussing different size options with the client.



Showing the client visualisations of the sofa using different fabrics.

STEP 5

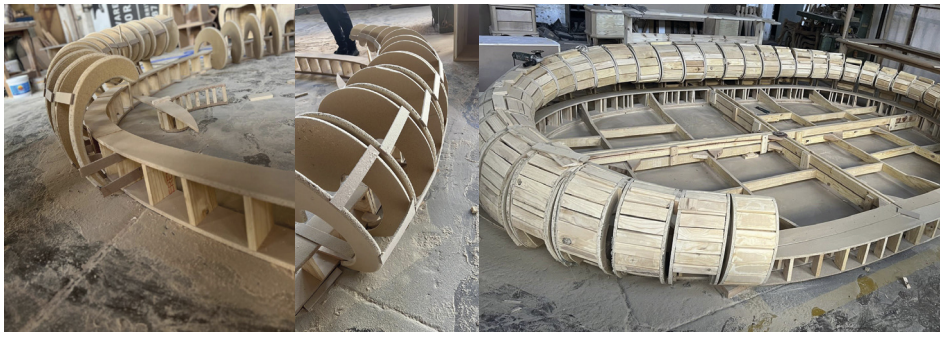
Collaborating with a manufacturer
Finally, connect with a manufacturer interested in producing the sofa. You can establish connections with manufacturers through Instagram or other social media, who may even commission you to design an item or collection for them. Engage in discussions regarding the design details, materials and manufacturing process.

STEP 6

Manufacturing the sofa
The production process will likely involve constructing the frame of the sofa according to the 3D model, involving precise measurements and attention to detail to ensure structural integrity. Due to the size, transportation can be a challenge, which may require the frame to be produced in separate pieces that can be reconnected seamlessly.

Once the frame is complete, the upholstery process can begin with the chosen material being carefully applied to the frame. Unique challenges may arise during the upholstery process, as each arm of the sofa may need to be upholstered separately and tied together, due to height variations required to follow the design concept. Challenges such as these require great skill and craftsmanship to ensure a seamless and uniform appearance.

The final step of the production process is the creation of a 3.5-x-2.5m custom, mattress-like cushion to go on top of the framework, ensuring comfort and removing visibility of the joints.



Production process - building the framework



Production process - upholstering

STEP 7

Final product and delivery
Upon completion, arrange delivery. The sofa may require international shipping, but marks successful completion of the project, from a Midjourney prompt to a comfortable recluse providing daily comfort to a happy client.



Final product.



Sofa delivered on site.

TIPS

Embrace AI as a creative partner, not just a tool	Approach AI with an open mind — let it surprise you. Use it to break creative blocks, explore unexpected ideas and push beyond your habitual design patterns.
Create beyond your own imagination	The best designs often come from surrendering personal biases and allowing something bigger to emerge. AI can act as an extension of your creativity, helping you explore beyond your own instincts.
Design for people, not just aesthetics	AI can generate stunning visuals, but beauty alone is not enough. Always bring the focus back to how a space will be experienced - emotionally, socially and functionally. Technology should serve the goal of creating meaningful, inspiring and livable spaces.

What does the future role of the architect or designer look like?

I believe the future role of architects and designers will be to shape how technology enhances the human experience in the built environment. While AI will excel in analysis and optimisation, our strength will lie in asking the right questions, making ethical choices and ensuring that design serves a purpose beyond functionality. Architects will focus on shaping broader visions, translating intangible human values into spatial experiences and fostering connections between people, nature and technology. Our role will not be to compete with AI in precision but to guide how technology integrates into human environments with empathy, emotional intelligence and a deep understanding of how spaces impact lives and societies.









The Beginning of the AI Revolution in Architecture

Tim Fu

DOI: 10.4324/9781003726500-8

‘A house designed by Zaha Hadid’. This description for an AI-generated image post that I found while browsing social media caught my attention.

Proclaiming to be a ‘completely synthetic AI-generated image’, this post featured an image of a curvy, blobby building that was a cross between an abstract interpretation of Zaha Hadid’s architecture and melted cheese. As someone who was working at Zaha Hadid Architects (ZHA) at the time, this piqued my interest. On the one hand, I was amazed by the technological marvel of an image wholly synthesised by machine; on the other hand, the the AI versions of Zaha Hadid designs were sub-par to the quality of design that we strive for at ZHA. It was for that reason I decided to try it myself.

This was the subtle yet crucial turning point of my career. For the next two years I was trying to push the task of designing in the style of Zaha Hadid using AI to its limit with methodically refined AI prompting techniques. I had gradually amassed a following during the AI boom when I showcased designs of my AI creations using my own artistic input and technical understanding. Later came interviews, lectures, conferences and exhibitions. By the end of the two years, my Instagram reach alone had amounted to millions. Interestingly, it is here that clients began approaching me in request of design services. I decided to leave ZHA to found my own office, Studio Tim Fu (STF). With a team assembled in London, we now continue to experiment and push AI in the design process with the aim of developing a practice that is uniquely organised to gain advantage of this AI revolution.



Tim Fu presenting his curvilinear visuals.

First and foremost, whether we love or hate the arrival of AI, we must all agree on one thing: it is disruptive to all creative industries. For architects, our values can be roughly summarised as a dichotomy of creative skill sets and technical skill sets. AI has shown varying capabilities in both directions. On the creative side, we can use diffusion text-to-image AI models like Midjourney and conjure up photorealistic and detailed architectural renders merely by describing our concepts with words. On the technical side, we can let AI amass large sums of quantifiable data to search for optimal architectural solutions. At my practice in London, I have encouraged my team to continue pushing the extents of AI in the search for practical use cases in the architectural design process. From our research efforts, we’ve discovered and adapted many AI-driven processes into our day-to-day practice. Organising my own design team has also a llowed me to explore alternative office synergies between the human and the machine, helping us explore new opportunities brought on by AI.



Retail detail. Midjourney.

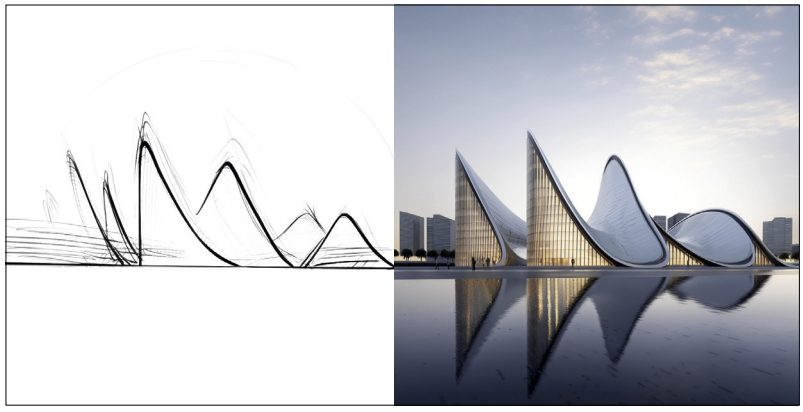
A Revolution in Creative Ideation

The diffusion AI model presents itself as an opportunity to reconfigure how we ideate. When we are given a project in a pre-concept/concept phase, we research precedence, both in traditional heritage and existing contemporary solutions, to spark our creative process and inspire us to come up with new solutions. Today, we can use AI as part of that conceptualisation process. Instead of relying on Pinterest to see the limited precedents for an 'Aztec-inspired seaside villa', we can describe it in Midjourney and see synthesised visualisations based on similar existing image datasets. In addition, we can describe it in a much more detailed manner to get much closer to what we need. We can then refine our results iteratively, giving us control of the outcome. This is what we at STF call 'prompt-crafting', which is a process that is undoubtedly elevated only with human artistry and technical experience. Our team now organises our design efforts on the company discord server and each member can prompt on Midjourney in their own channels. At the end of every exploration session, we 'pin up' our collection of ideas for discussion by sharing the best discoveries to the main channel. With a full team each sifting through and refining hundreds of ideas in their own prompts, we always end up with a large pool of high-quality architectural directions. With my professional background and visual acumen, I make the final decisions based on the pool of results, not unlike any other design directors.



Curvy villas. Midjourney.

ControlNet AIs are another front which revolutionised the way we ideate. With programs like Stable Diffusion, LookX or Rendair, to name a few, we can render designs from an image input. This has immense architectural applications on multiple fronts. We can develop accurate program massings in 3D and take a screenshot to be converted into architectural renderings. We can also rethink the process of architecture, by linking physical drawing processes to AI. For example, we can uphold the tradition of sketching by hand and then prompt it into a finished architecture. We can also bring back the childhood joy of playing with physical blocks and see how this develops into a commercial high-rise. I've also questioned our architectural process by crumpling a piece of paper and then converting that into architecture in the style of Frank Gehry or Daniel Libeskind. I don't think there is necessarily a right or wrong way of approaching the architectural design process and AI has only expanded the possibilities of how we can explore.



Sketch-to-render. LookX.



Playful blocks. LookX



Crumpled paper architecture. LookX.

A Revolution in Efficiency

Traditionally we were presented with a problem and given time to come up with a solution. This solution-searching process used to be manual and time-consuming. Our ideas were discussed in words, sketched by hand, modelled in 3D, iterated over and over in 3D space and then rendered as final imagery. We would then make visual judgements based on the render and then repeat this process to refine the design. Today, we can craft prompts and conjure up hundreds of image results with a degree of realism and finesse. This allows us to make a myriad of macro decisions based on form expression, materiality, lighting and other architectural factors. Simple decisions such as material and lighting combinations can make huge differences in the overall look and feel of a design. Today we can explore many combinations of these decisions and see in mere seconds the global effect they have. I believe this is partly where the revolution lies.

We can produce so many results at such a fast rate that we shift our focus from quantity to quality. We are going into an age where we are inundated with information and ideas. I like to call the AI age the 'era of abundance'. It is here that we will have to learn the art of distilling ideas from abundance. The act of being an architect will be about learning how to navigate a high quantity search space to distill the optimal results. Thus, a good architect will never be satisfied with the results of the AI but will learn to push the machine in the direction of their will.

While an abundance of ideas will speed up our ideation, huge efficiencies are also obtained in the visualisation of architectural ideas. One example is the need to render a developed architectural model. This process was either done in-house or through a professional rendering team and was always a time-consuming and resource-intensive endeavour. Traditionally a design team would always have to wait for the rendering team before giving additional feedback, so there is a delay in the workflow. Ray tracing is a graphics processing unit (GPU)-intensive process that eats up a lot of time, especially when things need constant fixing and updating. Diffusion AI is a non-GPU-intensive alternative that does not require modelled or digitised material parameters. Using predictive deep learning algorithms, AI renders results based on what it is trained on and the style references given – this process greatly reduces render production time. With AI technology developing so fast, we are beginning to see signs that architects will develop their own architectural AI models that can review design decisions in real time. This would connect concept design and technical delivery, greatly improving architectural design quality and performance.

A New Way to Run a Practice

Facing such disruptive technology like AI, I firmly believe that how we practice architecture will be turned on its head. Firstly, we humans have a limited learning capacity, which means we specialise in each phase of architecture/construction, and we partition our practice based on these specialisations. An architect may not have the best interest of the landscape designer, which may not have the best interest of the structural engineer, etc. The process of working together breaks the project into partitions that often creates inefficiencies and worse, incompatibilities. If you are an architect, I'm sure you've experienced countless times when your designs were negatively altered by an undesirable structural proposal, or sudden cost-cutting decisions. These decisions always happen in different stages, which is why a design evolves and devolves back and forth, inflating the time and cost of designing. With AI, we are approaching a super intelligence that can grasp all the desired outcome at each stage, and eventually provide a balanced guidance to each of the human specialisms. Our team is currently building an AI design system that connects to building information models (BIM) and structural performance data. It is still in its infancy, but we believe this is the direction in which things are heading. Overall, I believe that the pipeline from design to construction will be streamlined in real time and humans can best use their time to deepen their specialisms at each stage, like a highly skilled maintenance worker making sure that the pipeline of data runs smoothly and accurately.

Running an architectural practice in the age of AI, I have come to speculate that our future currency in design is data and compute. When it comes to data, everyone seems

to be on an equal footing, given we all have access to the same public visual data and can train AIs based on existing designs. Given this democratised starting point, having a proprietary database is crucial for every practice, so they can develop consistency in their solutions.

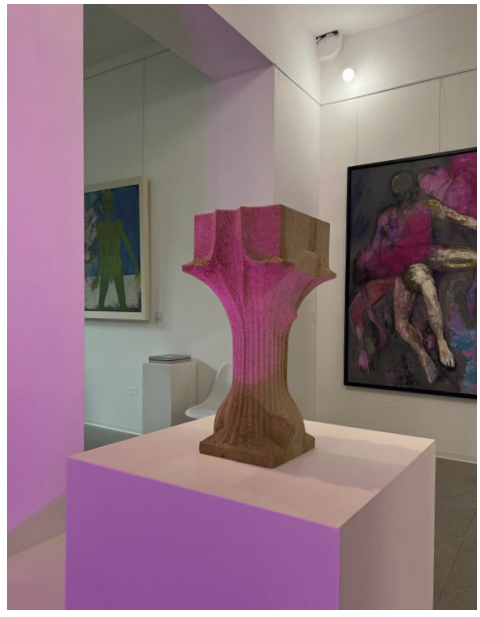
When it comes to compute, we are primarily using cloud services, each with their limitations of computational power. For example, when we want to develop more details in our design, we use AI upscalers. This program allows us to further add architectural details in our design. However, we need to zoom in and iterate multiple times at a higher compute to generate more detailed images. The accuracy and quality of our designs rely heavily on the level of compute that is available.



AI design-to-construction. Midjourney.



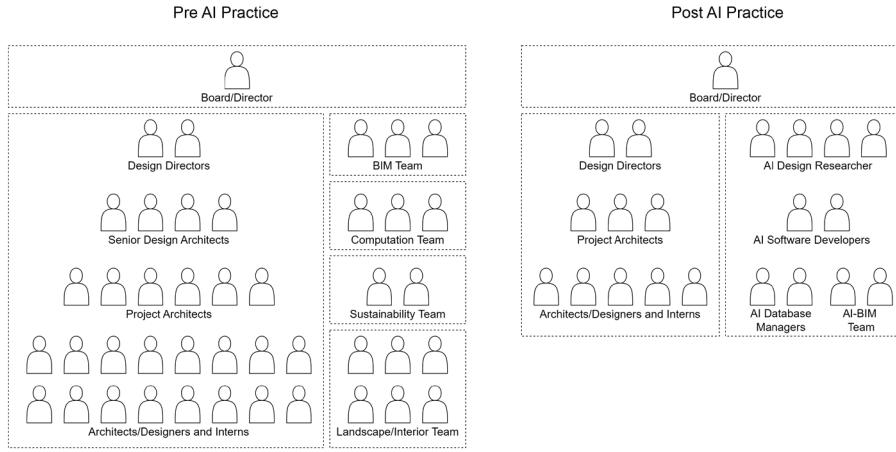




AI Stone Column - a collaboration with Tim Fu and Till Apfel.

A Future to Hope For

It is always difficult to speculate what the future holds in the midst of such rapid technological progress. For the architectural profession, there is bound to be organisational restructuring. From what I can see, as AI becomes more integral to design practice, we will dedicate more specialists to the maintenance and governance of AI. An architectural practice of the future could be a 50/50 division between a design team and an AI management team. Design teams will still be led by architects but the AI team will support all operations with potential new job titles such as AI research team, AI dataset managers, AI configurators and model-trainers, AI BIM managers, AI structural managers, AI operations, AI legal and finance team, and potentially many more.



The pre-AI vs post-AI practice.

AI is surely going to replace more of our tasks as we move into the future and we are left to ask: What is our role as human designers? I see it as more akin to driving a car with autopilot on. You are still the executive driver but you can set the GPS to tell the car where to go. You must still rely on your human instinct, intuition and knowledge, set the best destination and not allow it to veer off course.

Finally, we may have to face the hard-pill-to-swallow sooner or later: that AI will in fact reduce our workforce. By then, we can hope that enough AI advancement will provide us with universal income, and we can prioritise deeply human activities such as the creation of art and culture. We may favour architecture that is human made by the 'artisan' as opposed to industrial AI processes. The 'AI Stone Column' is a project that I created to exemplify this hybrid artisan future. Here, a stone column design is conceived with AI, but it is then digitally modelled in 3D by a human with parametric software, and then given to a masterful stone mason to produce by hand. An AI-driven process can still welcome human handicraft. I believe we humans won't always choose what is the most efficient or functional, but rather what is the most human and artful. A future where machines don't replace human art but elevates it is the future I am looking forward to.



Breeding New Typologies

Midjourney + Kohya LoRA Training

Will Garner

DOI: 10.4324/9781003726500-9

JOURNEY

The ‘Breeding New Typologies’ workflow harnesses the creative power of AI to redefine architectural styles by developing specialised AI ‘brains’ or fine-tuned models. Utilising a suite of generative AI tools, including Midjourney, Stable Diffusion and ControlNet, this method merges theoretical design elements with tangible geometrical forms. The workflow facilitates a smooth transition from abstract concepts to executable designs.

By incorporating a low-rank adaptation (LoRA) training method, the workflow enables us to train a fine-tuned model with specialised focus on creating outcomes similar to any set of images you upload and tag in your dataset.

Key benefits include the ability to incorporate and directly reference your own bespoke dataset while prompting novel variations. By mixing ingredients from any set of images that capture the concepts, aesthetics and site-specific elements of your projects, one can achieve image generation outcomes that truly capture the essence of a given look and feel.

Please note: using synthetic AI-generated outcomes as data for a training dataset is known as recursive training and could lead to ‘model collapse’. Although it is not recommended for open-source sharing and deployment, for personal use it can be a quick and effective way to create a bespoke, fine-tuned model if you use high-fidelity outcomes.

WORKFLOW

STEP 1

Exploring in Midjourney

The process begins by collecting our ‘ingredients’ – images and descriptive texts that capture essential aspects like contemporary apartments, materiality, form and shape. In Midjourney, these elements are carefully merged by adjusting the balance between image and text weights and introducing a chaos (--c) factor to spark creativity. This method facilitates the discovery and creation of new façade typologies, seamlessly integrating traditional architectural elements with contemporary influences, allowing a reinterpretation into new, unseen typologies and styles.



Gathering ingredients - combining images and descriptive text for initial input.



Results from Midjourney - creative exploration output.

STEP 2

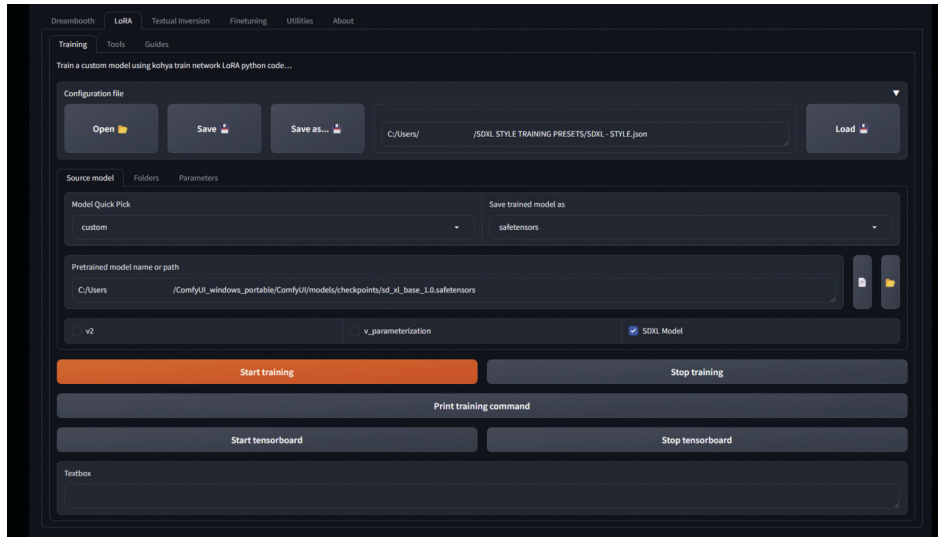
Establish dataset
From the exploration in Midjourney, a dataset of 15 to 25 diverse, high-quality images is created. The breadth and quality of this dataset are crucial, as it forms the foundation for the model's learning process, influencing the diversity and calibre of the architectural designs produced. An essential part of this process involves upscaling each image to ensure high quality, while also maintaining a manageable size to prevent slowing down the training phase. This balance is key to optimising the effectiveness and efficiency of the model development.



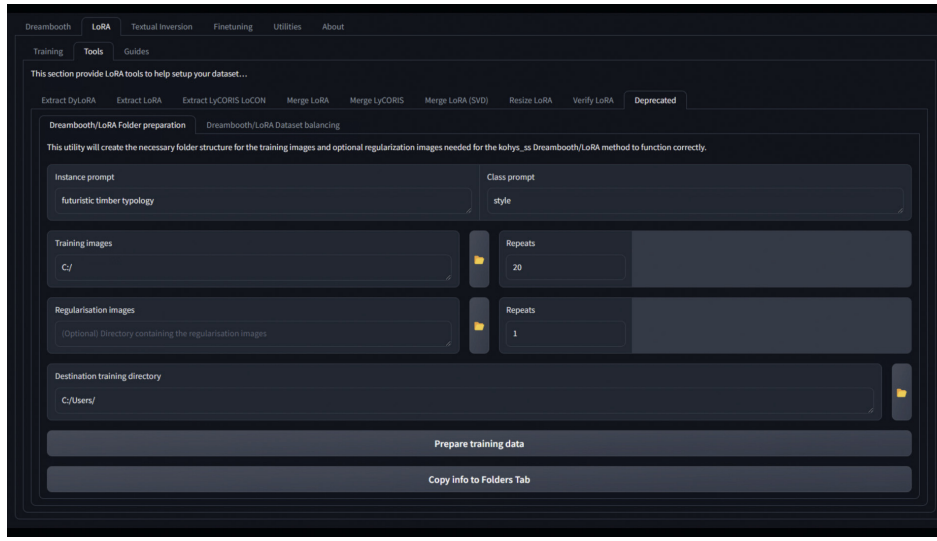
Establishing data for LoRA training – compilation of the dataset.

STEP 3

Labelling the dataset for fine-tuning
The next step involves meticulously labelling the dataset – a critical task for the effective fine-tuning of the LoRA model. Each image should be annotated with precise labels that capture key aspects such as architectural elements, materiality and overall scene composition. These labels are vital for guiding the model during training, ensuring it accurately replicates these design characteristics. You can streamline this process by using a tool like ChatGPT, which can automatically generate consistent, contextually relevant labels. Proper labelling is crucial in shaping the model's ability to learn and reproduce the desired architectural style.



Selecting the source model for LoRA training through Dreambooth.

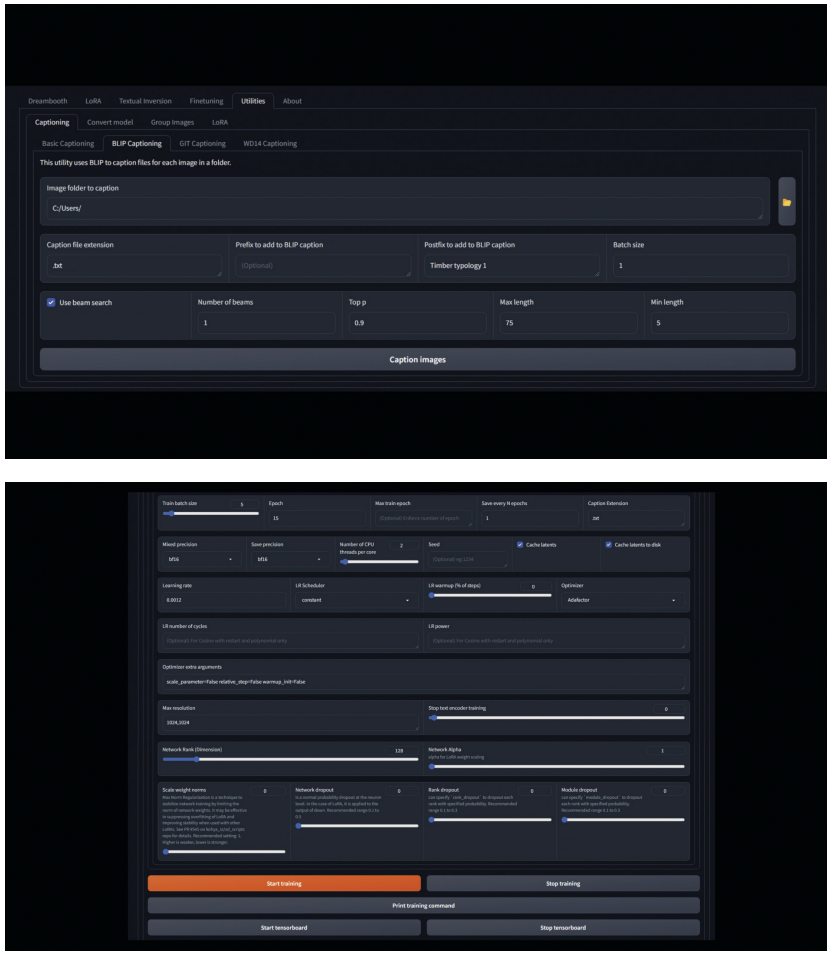


Directing the program to the image training files for input.

STEP 4

Fine-tune LoRA with Kohya

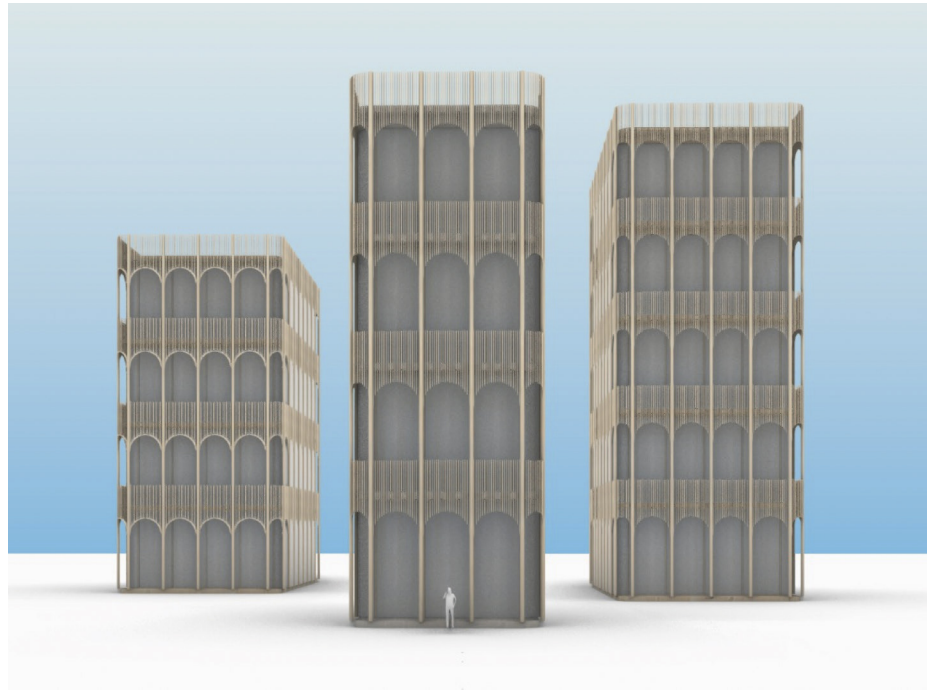
With the dataset labelled, the next step is to fine-tune the LoRA model using Kohya. This involves adjusting key parameters like learning rate, batch size and the number of epochs. For example, increasing the learning rate can help the model adapt faster, while more epochs allow it to refine its understanding through repeated passes over the data. Experimentation is needed to be able to effectively tune these settings in a way that ensures the model captures the architectural nuances in the dataset while remaining flexible enough to allow for creative exploration and adaptation throughout the design process. More detailed Kohya LoRA training guidance can be found online.



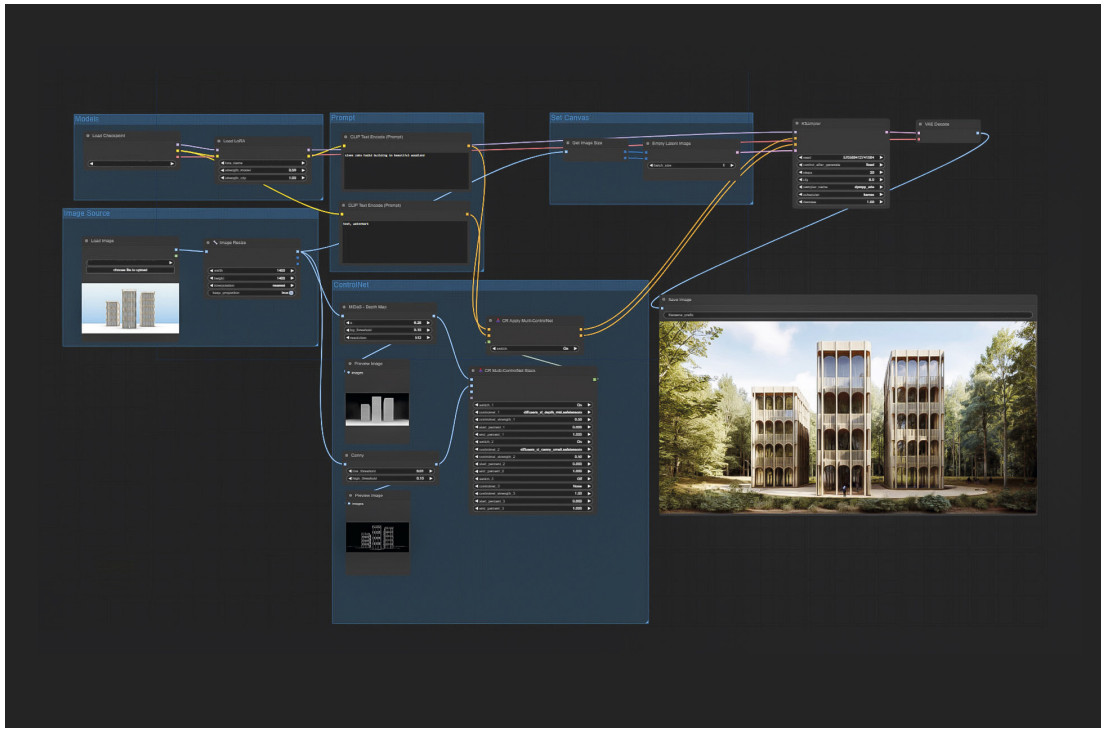
(Top) Captioning images for training through BLIP.
(Bottom) Configuring training parameters.

STEP 5

Render with ComfyUI and Stable Diffusion
In the final step, the fine-tuned LoRA model is deployed through ComfyUI and Stable Diffusion to render new images or enhance existing geometries, such as 3D models and sketches, using ControlNet. This flexible approach allows for the creation of novel architectural designs within the established style or the adaptation of pre-existing data into cohesive visual outputs. By leveraging the LoRA model, architects can gain more control and seamlessly integrate AI-generated concepts with practical design elements, effectively 'breeding' new architectural typologies. This process ensures a continuous evolution of design languages as well as the varied reproduction of existing ones.



Capturing modelled massing for LoRA rendering.



ComfyUI in operation - rendering developed geometries.



Example outputs and further visualisation results.

TIPS

Know when to use AI	Understand that AI is a powerful tool, but not a replacement for the foundational principles of good design.
Stay grounded in your subject	Strong design fundamentals will always matter, even as technology evolves.
Keep making	There's no better time than now to experiment, build and innovate, especially as new tools and capabilities emerge.

What does the future role of the architect or designer look like?

I envision the future role of architects/ designers as the custodians of the creative process, seamlessly collaborating with diverse AI systems - and eventually robotic counterparts - on what will be quite the journey.



The Effects of AI on Practice

Keir Regan-Alexander

DOI: 10.4324/9781003726500-10

Before GPT-3 and After

During the changing seasons of 2022, three easy-to-use AI models were released in quick succession, seedlings that brought change for designers of all kinds: Midjourney v1 in Spring, Stable Diffusion v1 in Summer, and in late Autumn of the same year, Open AI's GPT-3. Since then, progress has been relentless, with model parameter size expanding on a logarithmic scale.

With increased model size we see broader possibilities for real-world use cases. A primitive large language model (LLM) can produce novel rhyming couplets, whereas a large parameter model such as Claude 3.5 or OpenAI o1 (with many hundred billions of parameters) can now execute demanding tasks such as debugging computer code, assessing financial strategy and analysing legal risk.

Humans, whose knowledge is constructed mostly from interactions with linear phenomena, tend to think in logical sequences; something begins, steps are followed and an expected outcome is reached upon completion. But when model parameter size is multiplying by ~10X with each development cycle, we can become quickly bewildered by the speed of exponential change. The full capabilities of each model release are only truly uncovered as users and developers test their limits.

The open-source culture of AI causes a superfast product development cycle. When this speed is combined with the abundant access to shared information online, the outcome has been a thrilling hype cycle that keeps rolling. Just as we become accustomed to one breakthrough, we are immediately presented with several more.

In future decades, I anticipate we will look back at creativity in the design fields under two main epochs, before generative pre-trained transformer (GPT) and after. It is not that I believe our creative purpose will be fundamentally replaced by AI, but more likely that the two become so intertwined and interdependent as to be indistinguishable from one another in a post-GPT world.

In such a creative landscape, huge financial opportunity and moral risk will abound and the stakes will be high for the field of architecture. It's not hyperbolic to state that the physical environment of our future cities will be moulded largely by the work of computer scientists and the algorithms they write in the next few years.

Just Another Tool?

Many people explain away AI as 'just another tool' to be filed next to set squares and AutoCAD – a utilitarian invention that will aid productivity. Yes, it is a tool, but it's the first tool that has its own synthetic judgements to make. A tool whose output is readily indistinguishable from human expression.

To label AI as simply a tool is to downplay both the opportunities and risks it brings. If you decide to integrate AI and care about your creative process, then you should be aware that in time it will gradually alter the way you work.

When we call something a tool, we presume to be the master of it, and while this explanation may feel comforting, it lacks a deeper reflection. If we are to reach artificial general intelligence (AGI) as is OpenAI's stated mission, then it's very possible that we won't perceive when the tool has become the master of us. Scale this big and

(Previous page) A sample board inspired by the colours of an interior scene render for a museum in London, produced using text and image prompting in Midjourney.

scary idea down to your own creative process within your practice and the same can apply – at some point, an AI you are using will be influencing your decision making and you may not be aware in what way.

Jeff Bezos describes how we are co-evolving with our tools and how this process can lead to unintended consequences: 'We co-evolve with our tools. We invent new tools and then our tools change us… for example, we carry around in our pocket a smartphone, and one of the things that phone does is act as an attention-shortening device, because most of the things we do on our phone shorten our attention spans. That's one of the ways we're co-evolving with that tool.'

The very same could happen with creativity in our field.

A recent study published in the journal *PNAS Nexus* has introduced the new term 'Generative Synesthesia', a label coined to describe the enhanced state of creativity made possible by working with an AI co-creator. The study's conclusions suggest that the outputs from creators using AI were perceived as more favourable by viewers over time and creators also reported consistently higher levels of productivity. These findings align exactly with the value proposition being sold by AI companies.

But something else also happened that is significant for creatives. For the broad base of users, an overall convergence of styles and tastes occurred; the average novelty from image to image measurably declined. While more output was being created to a higher standard, everything was becoming similar. At the same time, so-called 'peak novelty', while rarer, was more pronounced from a smaller number of users.

Research remains scarce for now and definitive conclusions are not yet possible, but the essence of these findings instinctively rings true when compared to my observations watching early adopters work in practice. The AI tide may raise all participating boats somewhat, but there will likely remain creative outliers who capture the biggest beneficial gains because of their unique insight and qualitative judgements that allow them to stand out in a saturated field of competition.

Is it All AI?

To create architecture, we need to balance precision and nuance, like two complimentary hemispheres of the same brain. An architect who is only good in one hemisphere is unlikely to consistently succeed in the profession – you have to be able to do both.

On the precision side, the technology broadly referred to as 'generative design' is relevant. Software in the generative design realm focuses on the quantifiable. These programs ruthlessly execute procedures based on the deterministic parameters (rules) defined at the outset. They give 'predictable' results that produce vectors with coordinates that can be measured in millimetres, ideal for the prescriptive realm of construction-ready design.

On the other hand, software that uses generative AI technology applies to the qualitative aspects of practice and tends to sit more in the fuzzy realm of image creation, immeasurable pixel rendering and copywriting. This tooling is more unpredictable and is best suited to assisting specifically in creative workflows.

I have seen much confusion in practice about what is a generative AI technique and what is a parametric technique because, for marketing reasons, people are often using these terms interchangeably. A quarter of global venture capital deployed between 2023–4 has gone into generative AI products, a doubling on the previous year.

A trend that is set to increase – so no wonder everything is being labelled as AI, whether the title is accurate or not.

Many of the most fervent early adopters hail from the parametric tradition and are using both deterministic generative design tools (like Grasshopper) in combination with generative AI tools using a node-based interface for Stable Diffusion, such as ComfyUI. To the uninformed, the outcomes will appear like magic, fluid geometries extruded, twisted and honed while being rendered in real time.

Having these two oppositional definitions as a framework for two sides of a wide spectrum should help you analyse and navigate any new technique you observe that appears at first inexplicable.

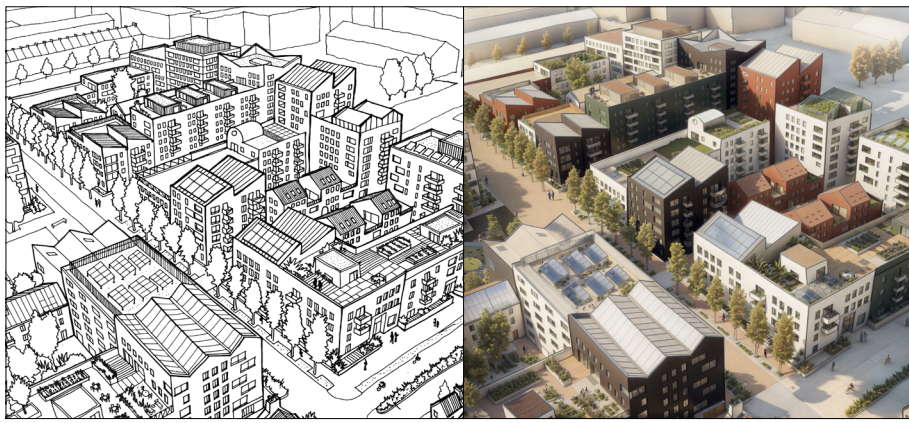
Faster, Yes, but What About Better?

A widely held belief is that generative design and AI will make things faster and all the early evidence tends to prove this to a dramatic extent. Seldom will you see an example of a new technique without the vanishingly small window of time that was taken to create it being presented alongside.

So, will these innovations make things quicker? Yes, but I think there is a more interesting conclusion – perhaps we change the way we solve problems instead. We may be able to do it quicker, but can we do it better?

We measure the known unknowns on a project with the term 'risk'. Compliance risk, planning risk, construction risk, health and safety risks are all worries that can be quantified, weighted probabilistically into the business plan and gradually mitigated through analysis and a deeper level of early planning.

Being able to better predict and plan for things that won't work is a great way to bias the probability of success in your favour. If we can get a 'fairly decent' early idea about the project performance for say daylight, wind, sunlight, material appearance, development density, area efficiency and carbon use intensity then we can reduce the risk of designing something that will be a banana skin later.



A generative design and AI workflow from practice - a digital hand sketch over a massing geometry in Spacio, (a parametric urban planning tool).

A generative design and AI workflow from practice - a sketch-to-render produced in Stable Diffusion, using a LineArt ControlNet.

In tech, I have heard this technique referred to as 'premortem'. The intention is to expend some energy working out why something is a bad idea before you invest too much time into it. Another way to put it would be using design sprints to validate an idea for better or worse using shorter but more frequent design cycles. Thus, we keep bad ideas alive for the shortest amount of time.

If we can have a greater than 80% accurate answer a stage earlier and without great cost to the client then we should take those odds and find ways to exclude poorly performing design options sooner in the process.

Such a change would reduce the need to keep parallel (divergent) options live at the same time – something most architects would recognise as a contemporary problem in practice that reduces both profitability and morale among teams. Changing the way we organise ourselves to work in more concerted design sprints targeting the validation of an idea would mean getting answers more quickly, focusing our energies on quality and reducing the chances of a bad result later.

A generative design and AI workflow from practice – a digital hand sketch over a massing geometry created in Spacio (a parametric urban planning tool); a sketch-to-render produced in Stable Diffusion, using a LineArt ControlNet.

Customisation

Once you have looked at emerging generative design and AI products in detail, you may start to feel that all roads lead to customisation. This is because most fields of design have game theory at their core. The idea is that only a unique insight can unlock a problem or serve a need in a specific way and that projects succeed or fail on this basis.

In the built environment, the value proposition of a good developer or architect is that they can realise ideas for a brief that no one else can see or achieve, be it through a specific construction method, supply chain, development funding model, quality of design, specific planning strategy, etc. For this reason, I am sceptical that any multifunctional and generic development appraisal or assessment tool can ever truly succeed at scale without permitting heavy customisation by users.

Clients have never sought completely generic solutions or valued answers that are essentially available to anyone. Information and solutions that everyone can access present no asymmetry and therefore no opportunity. This is the likely outcome of using a generalised, closed-source 'black box' AI tool in practice and on projects. The solution? Don't use general tools, pick them up and customise them by pulling your own IP into them so that you can get answers that no one else can. In the near term, I see practices customising their approach in the following ways:

1. Text-based: via LLMs. For example, by preparing critical documents so that they can be used in repeatable workflows using long-context window 'Pre Made Prompt' formulas across departments.

2. Image-based: via Image Diffusion. For example, by creating custom style LoRAs, trained on in-house imagery and design methods. -

3. 3D-based: via generative design tools. For example, using workflows that are derived from in-house sector based template prototypes.



Customised image models. Early design stage concept for an HQ building, produced with a custom-trained LoRA for Stable Diffusion, created from synthetic training data, called ArkaPainter-XL.

These strategies all produce outputs that are 'local' to their practice, based on data that is privately held, unique (because no one else can get the same outputs) and the product of their own core thinking on previous projects.

A Changing Profession

The journey towards wider adoption of AI in architectural practice is as much about cultural adaptation as it is about technological integration.

Something I've observed working with AI transformation planning in practice is that a new way of doing something may exist and it may be superior, but creating the prerequisite behavioural change needed to realise those benefits is another, more human challenge altogether. Even as a practitioner who proactively strives to use AI-centric methods in my work every day, I still have to remind myself halfway through certain routine tasks that I could be using an AI to help me.

As graduates enter the profession in the next few years, they'll do so having learned about architecture through the lens of generative AI, perhaps by using tools like Midjourney in combination with model-making and sketching in order to test initial spatial concepts. To these future architects, AI will be perceived as a general-purpose technology that is already fully integrated into the way they think and solve problems.

AI is often compared to the invention of computer-aided design (CAD) or of building information modelling (BIM), but I think this is largely wrong. AI will most likely bring horizontal change to practices of architecture, by which I mean it will

in some way touch every department. In this way it is fundamentally different from CAD or BIM which are more vertically managed technologies.

Particularly when we consider the application of LLMs, it is useful to think of AI as a wide, horizontal umbrella covering all parts of the business: projects, support, bids, comms, etc. Yes, you will need technical expertise and support at the core of your approach, but most importantly you should establish design leaders in the practice who have the motivation and desire to try new techniques. These designers should be at the heart of the conversation and they should be the ones doing testing on live projects with technical support, in order to realise progress and decide what is useful to their design method (or not).

From what I've seen in practice, open-ended research projects can lead to naval gazing and indecision and won't realise actual progress. Much better to start with short, experimental bursts within real project environments where outputs are subject to rigid constraints and higher stakes. Establishing quick wins on real projects compounds over time and will bring more lasting behavioural change.

A recent survey in the *Architects' Journal* has suggested that in Q1 of 2024, half of respondents are already on some kind of adoption AI journey and the 'Early Adopters & Innovators' cohort are making up about 15% overall. This figure maps closely to the classical distribution model predicted in the diffusion of innovation theory by E. M. Rogers: a bell curve diagram that is applied to any new transformative technology.

In the coming years, if current trends continue, we will find ourselves operating in an environment where we are competing with a field of both AI-augmented architecture practices as well as AI-powered software products that also begin to offer parts of our scope of services to clients directly.

Many architects feel that the scope of their role has narrowed in recent decades, but by harnessing new tools that leverage our efforts and insights, we have an opportunity to actually expand the scope of both our responsibility and creativity – if we lean in and take judicious ownership of this new technology and put it to good use.

With such transformations afoot, the profession must also open up a critical dialogue about how we can maintain and reinforce our core value proposition into the future to encourage a more resilient mode of practice, one that can continue to be financially viable and that can deliver the highest quality of project output.


A
in
Q W E R T U I U Y O L
A S D F G H J K L L
Emter
Z X C V B N M Y M

AI Form-Morphing

Stable Diffusion Automatic1111 Prompt Travel with XL LoRAs

Ismail Seleit

DOI: 10.4324/9781003726500-11

JOURNEY

Creating an AI-generated animation involves a blend of creativity and precision, starting with the conception of using two specialised models: one for architectural forms (FormFinder) and another for architectural illustrations (Petals). These models undergo extensive training over one to two months, involving dataset preparation, captioning and refinement. The aim is to produce visually captivating, fluid animations that showcase intricate architectural designs. The process employs various tools, including prompts, seeds, classifier free guidance (CFG) scales, low-rank adaptation (LoRA) weights, samplers, steps, aspect ratios, styles and checkpoints, all meticulously fine-tuned. AI facilitates this workflow by acting as a creative partner, generating images based on text prompts. It doesn't dictate the process but rather collaborates with the user, who refines and guides the AI's output through parameter adjustments and selection of the best results. This interaction ensures that the final animation is both technically sophisticated and artistically engaging.

WORKFLOW

STEP 1

Model training
The workflow begins with an extensive process of model training that lasts for about a month. The models being trained are relatively small 'attachment' models called low-rank adaptation (LoRA) of large language models (LLMs), which is an efficient way of teaching diffusion models new concepts. The process involves preparation of a dataset of around 50 images and manually captioning each image with a focus on subject or specific elements on the image that you want the model to put more attention to. In this workflow two models were trained.

The first model called FormFinder-XL was trained to understand complex architectural geometric forms and was able to replicate these forms using different styles of images. The second model called Petals was trained to understand flat-shade architectural illustrations and was able to generate various images of this style using prompts.



FormFinder-XL model training and testing.



Petals model training and testing.

STEP 2

Model mixing

The next step involves an iterative dialogue between the designer and the diffusion model using different weights of the trained models. The main aim of this process is to find the best weight balance between the two trained models to get the desired aesthetics. In this case, the aim was for Petals to apply the illustrative style and FormFinder-XL to drive the geometry.

Here is an example of a prompt that involves both models with different weights: Architecture design illustration, hand drawn, clean lines, pastel colors, super minimal biomimicry sculpture structure, made of carbon nanotubes, form inspired by natural geometry, biomimicry, thick forest with wild grass environment: <lora:Petals:0.7> <lora:FormFinder-XL:0.5>



First attempts to mix both models together with different weights.

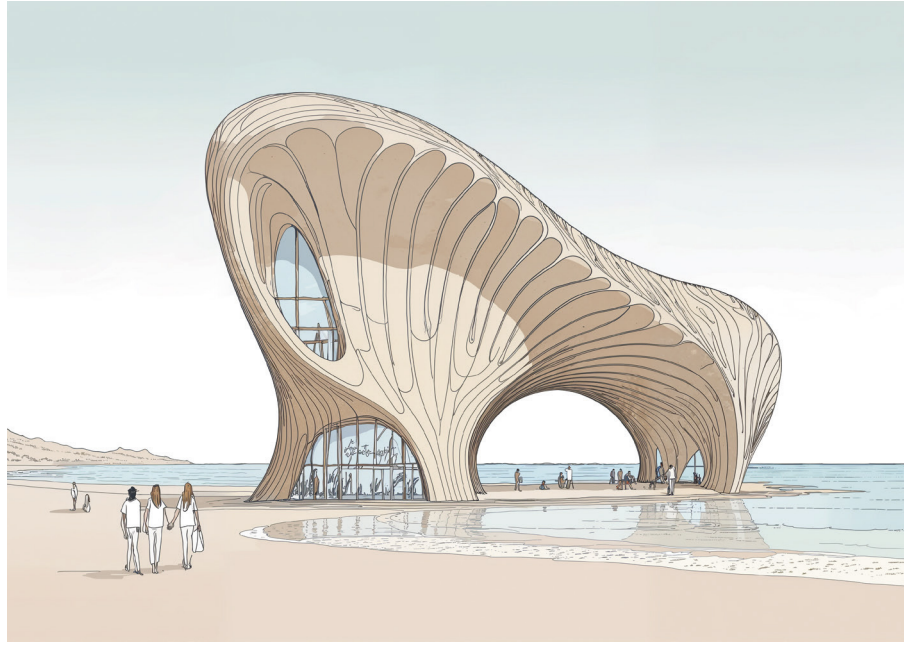
STEP 3

Setting the parameters
After a long process of discovery, an initial image has been selected. This image has very specific settings and generation parameters that are to be used as an initial input for the animation and form-morphing in the next step. Here are the generation parameters used to generate this specific image:

Prompt
Architecture design illustration, hand drawn, clean lines, pastel colors, super minimal biomorphic sculpture structure, made of timber, form inspired by natural geometry, biomimicry, serene empty Mediterranean beach environment <lora:Petals-XL-000002:0.7> <lora:FormFinder-XL:0.5>

Negative prompt
realistic, cgi, photography

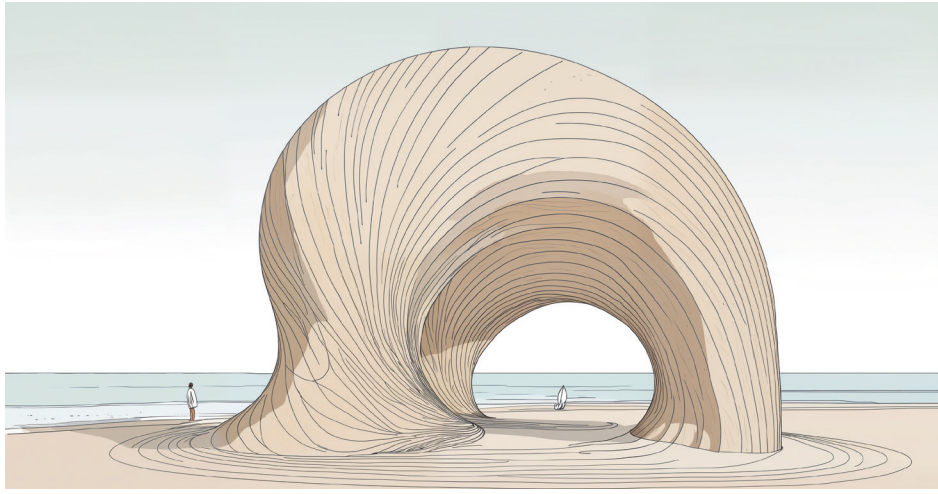
Parameters
Steps: 30, Sampler: DPM++ 3M SDE, Schedule type: Exponential,
CFG scale: 8, Seed: 3231938942, Size: 1536x1024, Model hash: 912c9dc74f,
Model: realvisxlV40_v40Bakedvae, Version: v1.9.3, UI: Automatic1111 -

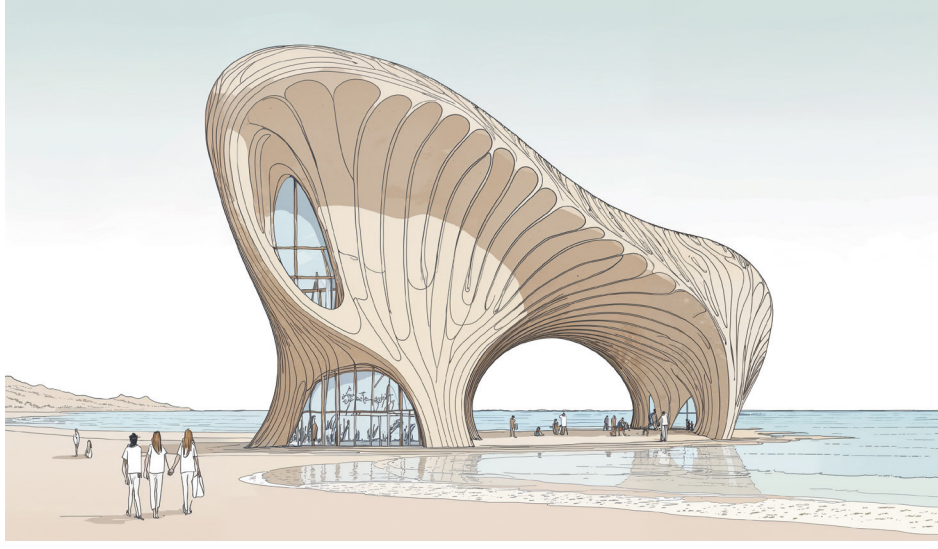


First frame image generated and selected with specific parameters and model hashes.

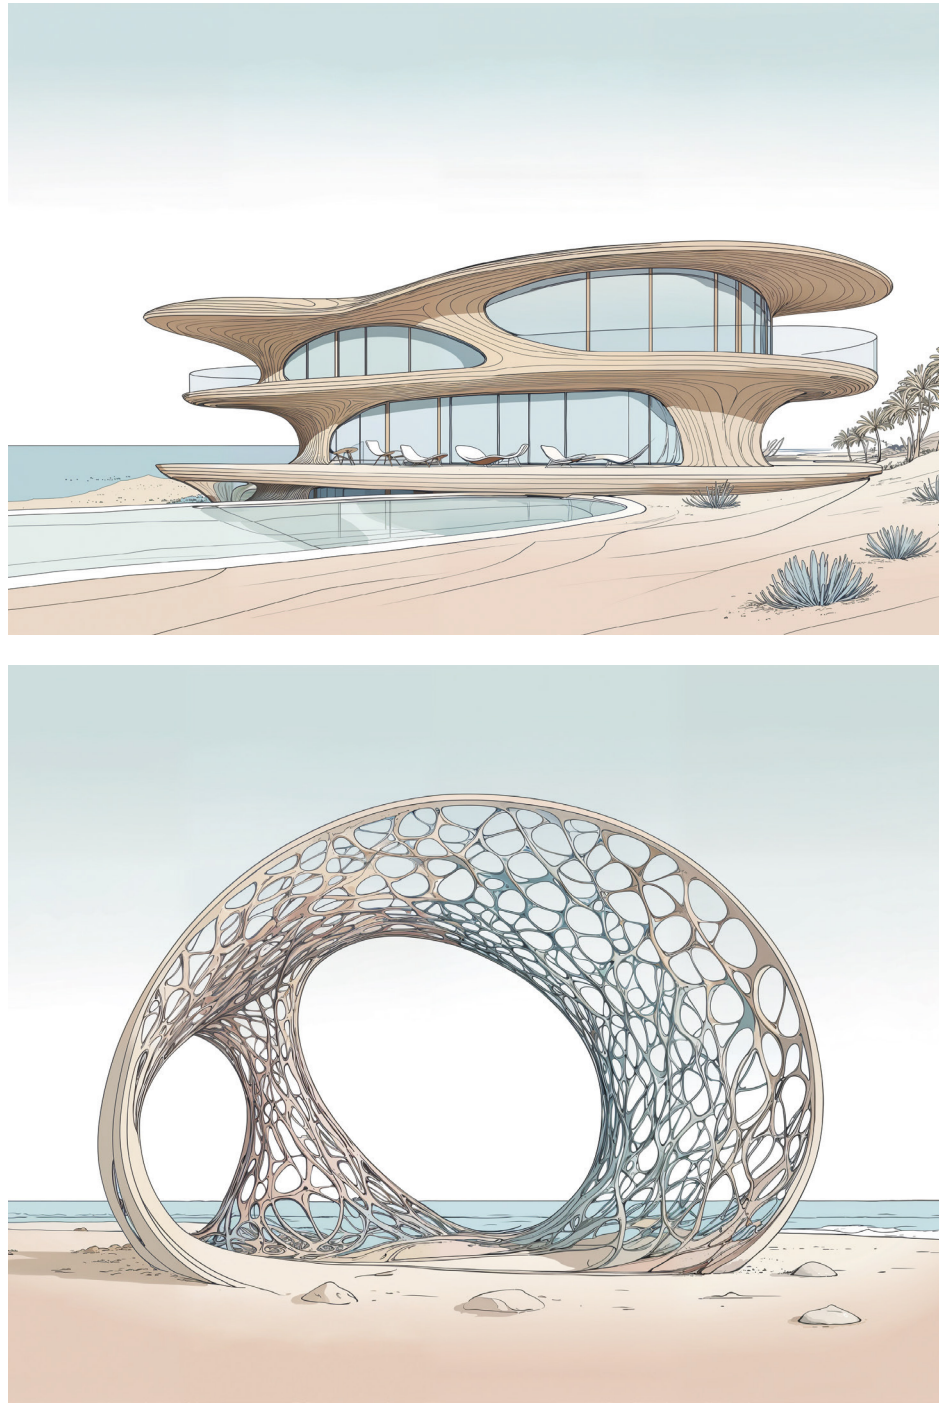
STEP 4

Prompt travelling
Using the same seed and settings specified in the previous step, four additional images that will act as keyframes have been generated and picked by only changing parts of the prompt. This leaves us with five images that will be added to a script that travels between them, generating in between frames. The settings used for this script are shown on page 116.

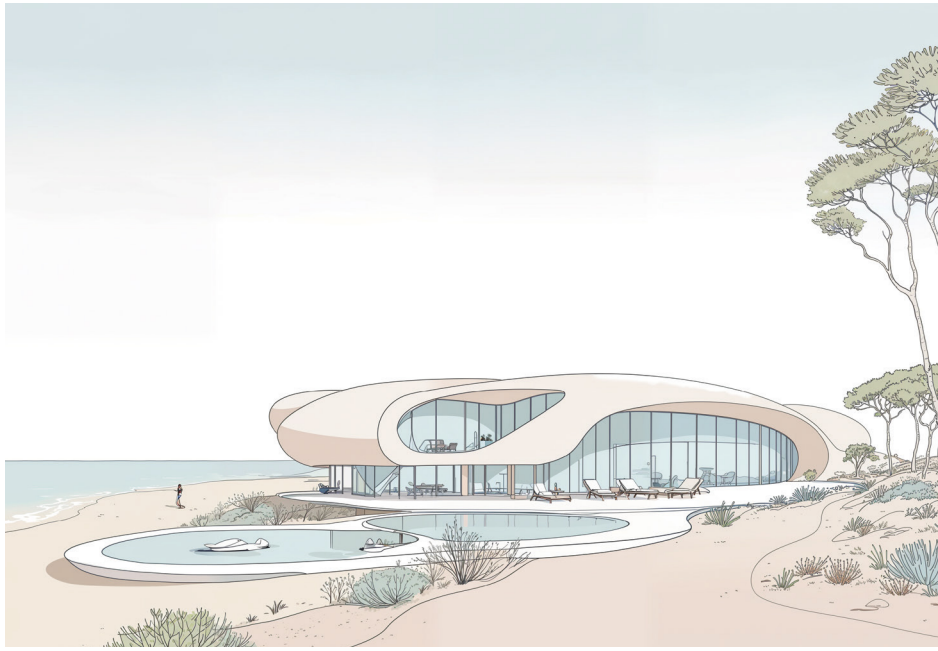




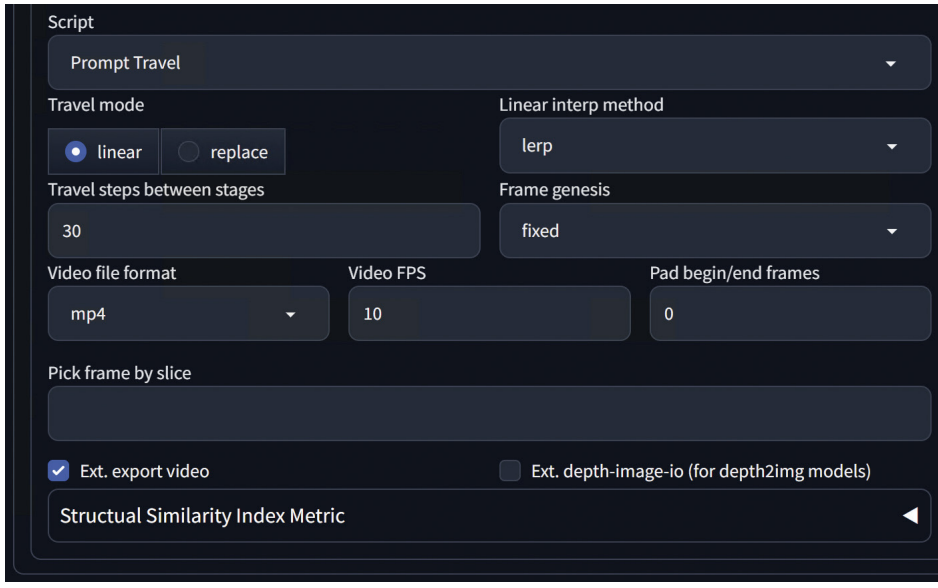
(Top) Frame 01 of animation.
(Bottom) Frame 61 of animation.



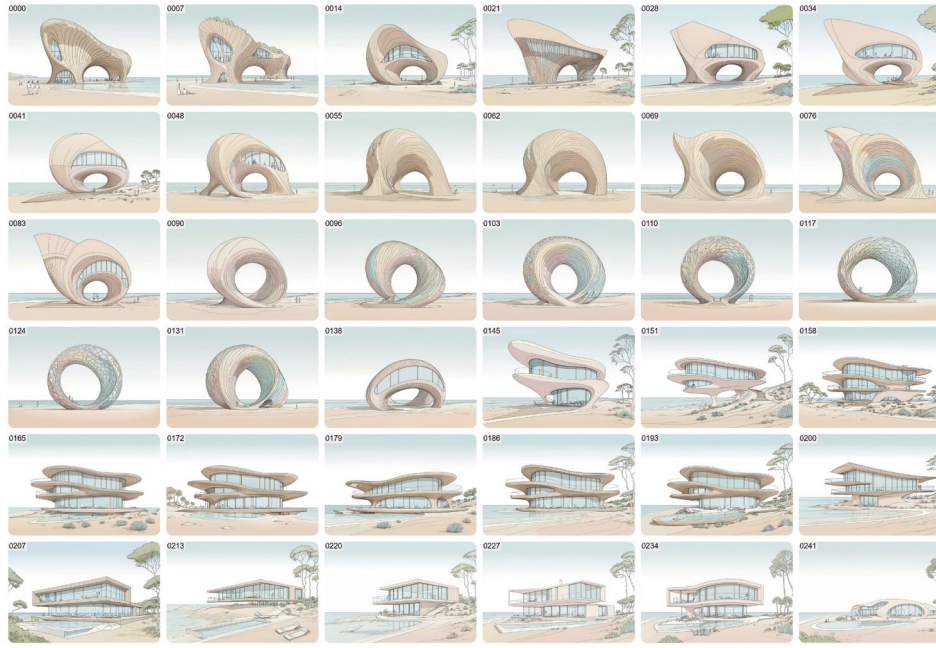
(Top) Frame 121 of animation.
(Bottom) Frame 181 of animation.



Frame 240 of animation (last frame).



Settings used inside Automatic1111 for Prompt Travel script downloaded from extensions tab.



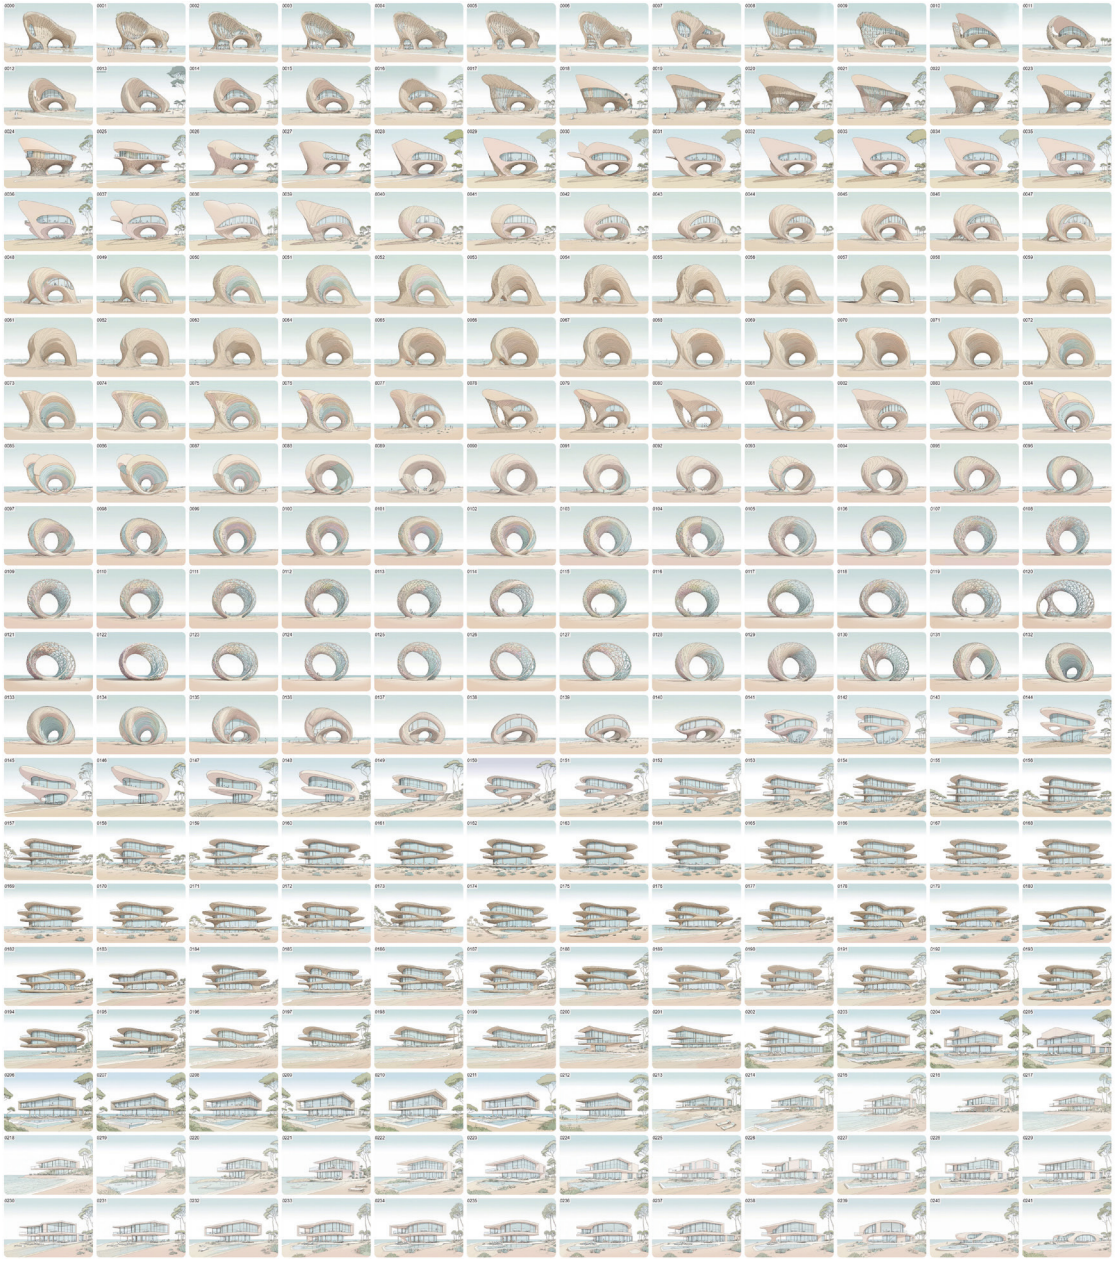
Excerpts of the generated 240 images with Prompt Travel.

STEP 5

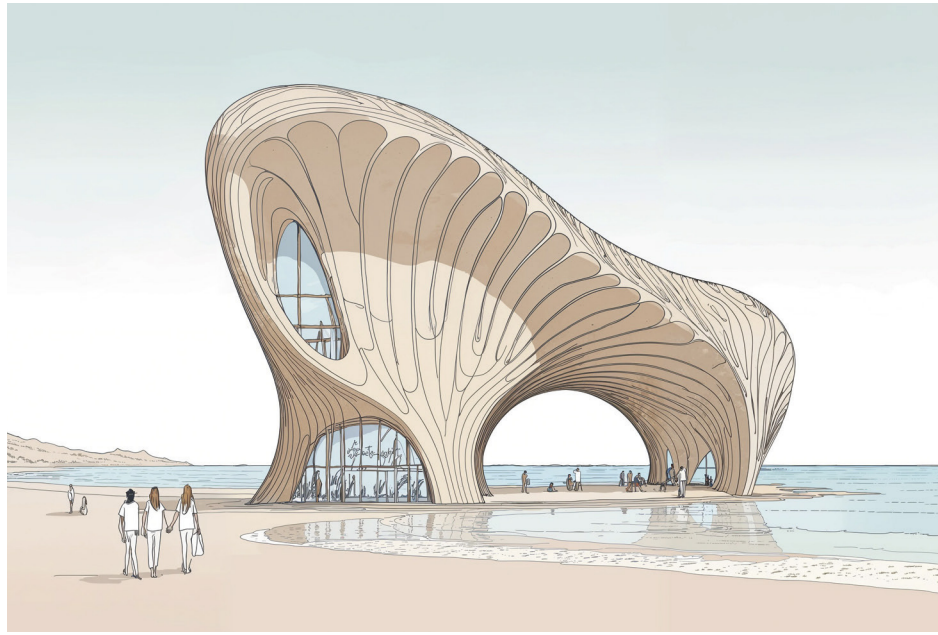
Upscaling and frame interpolation

Upscaling is then done to all frames generated in the previous step (240 images) from a resolution of 1536x1024px to 3072x2048px using ControlNet Tile model with SD upscale script and 4x-UltraSharp upscaler. All inside Automatic1111 webUI using Stable Diffusion XL model. The upscaling process not only increases the resolution of the images, but also adds up necessary details and cleans up artifacts in the process.

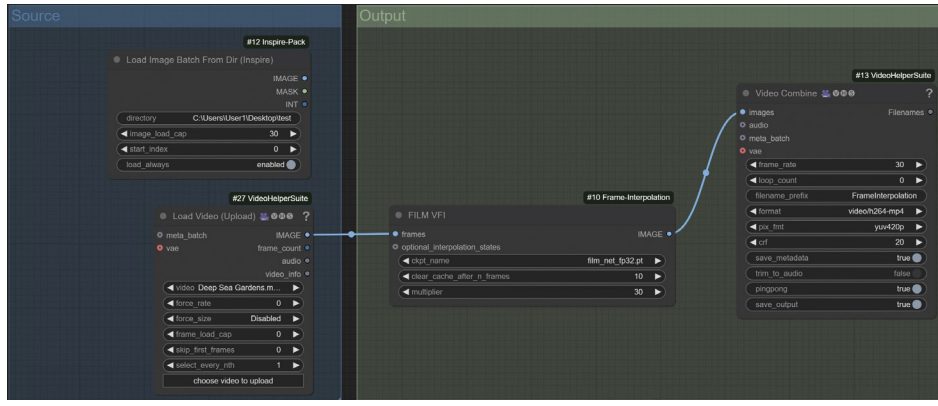
Frame interpolation is then undertaken; a process of inserting extra frames inbetween the frames generated to give a final, smooth, morphing animation. This is done using RIFE 4.7 frame interpolation AI model inside ComfyUI. The settings used for interpolation are shown on page 119.



All 240 upscaled image frames in the animation.



RAW generation versus upscaled image.



ComfyUI workflow used for frame interpolation, adding more frames inbetween generated frames.

Final Video

FPS: 30
Number of frames: 2880
Resolution: 3072x2048
Duration: 1 minute 36 seconds

TIPS

Master parameter tuning	Beyond prompts, settings like CFG scale, sampler and steps are as important as brushstrokes to a painter. Experiment with these to control and finesse the AI's creative range.
Be a curator, not just a creator	The key is knowing what to keep and what to discard. Generative AI offers endless possibilities, but your role as a designer is to select, refine and elevate what aligns with your vision.
Fine-tune your own models	The most impactful way to create distinct outputs is by training custom models, like LoRA fine-tuning. Tailoring models to your domain allows you to design outputs that truly resonate with your creative vision and aesthetic preferences.

What does the future role of the architect or designer look like?

I envision the future role of architects/ designers to be that of creators of bespoke generative systems, where they no longer just use tools but craft their own through processes like LoRA fine-tuning. By designing their own models and training them on curated datasets, architects will imbue AI with specific expertise - capturing the essence of materials, regional aesthetics or even cultural narratives. This ability to shape the DNA of their tools will enable designers to work with unprecedented precision and creativity, bridging the gap between abstract ideas and executable realities. Architects will become ‘model-makers’ in a digital sense, where their craft lies in defining the frameworks within which creativity unfolds. This new role ensures that human intuition, ethics and artistry remain central, using AI not to replicate but to redefine what is possible in design.



Architecture at 60fps

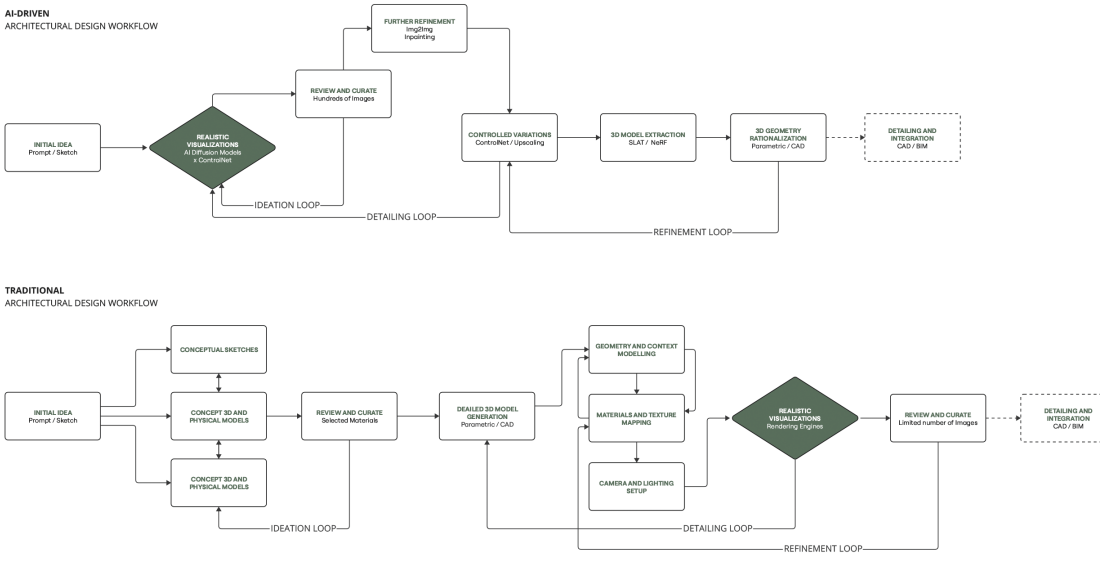
Carlos Bañón

DOI: 10.4324/9781003726500-12

Architecture has always been a dialogue between vision and reality, shaped by the tools we wield, from the pencil to the pixel. With the rise of AI, particularly diffusion models, this conversation has entered a new era. This paradigm shift is transforming the architectural creative process, unlocking unprecedented design capabilities while introducing new challenges and opportunities for human-AI collaboration.

The journey from concept to creation is inherently fluid, with deviations from the original vision being crucial to the creative process. Initial sketches ignite the first sparks of an idea, yet they also carry biases that can gradually steer the design off its original course. As the design takes shape through physical models, material constraints and scale limitations further influence its direction. Traditionally, achieving a convincing rendering required multiple steps – from crafting a detailed 3D model to refining materials, textures and lighting. These steps necessitated that the project be well-defined before it could be visualised in such detail.

However, recent advancements in AI technology are radically altering this traditional sequence. Diffusion models in particular reverse this process by generating highly realistic visualisations at the outset, without the need for a fully defined project. This approach eliminates many of the intermediate steps, allowing for rapid exploration and providing architects with unprecedented agility and potential for experimentation. We can now explore a multitude of design possibilities in the blink of an eye, where the density of information and the speed of data generation and curation have reached unprecedented levels.



Comparison between traditional and AI-driven architectural workflows.

This shift demands the creation of new AI-driven workflows, fundamentally challenging the traditional methods architects have relied on for decades. Behind the polished visualisations lie intricate layers of data and processes that no longer align with conventional architectural sequences. In this chapter, we will examine this new domain of design – architecture at 60fps – exploring the tensions and opportunities that emerge as we introduce elements that offer unique ways to communicate between human-produced material and AI systems.

Reconfiguring the Architectural Workflow

The integration of AI into architectural practice is reshaping the very foundations of the design process. Historically, the progression from concept to final design has been a linear and time-intensive journey. The sequential nature of traditional tools often constrains the exploration of alternative designs, as each revision demands significant effort and resources. For example, in a conventional workflow for a residential project, an architect might spend weeks iterating through hand sketches, computer-aided design (CAD) drawings and 3D models before arriving at a convincing design, with each revision cycle potentially taking days.

In contrast, the new hybrid human-AI paradigm inverts this sequence, offering a more agile and iterative approach. High-fidelity visualisations can now be generated immediately, bridging the gap between abstract concepts and concrete representations. AI's ability to rapidly translate text descriptions into visual forms further narrows the divide between ideas and their manifestations. This shift enables architects to refine multiple aspects of a project, such as materials, lighting and form, in parallel, without the constraints of a sequential process.

However, this transformation raises profound theoretical questions about authorship, control and the evolving role of the architect. As AI takes over more of the granular tasks of design development, architects must reconceptualise their methods to fully leverage its capabilities. Rather than detailing every aspect, we should focus on defining the critical elements that AI can then expand upon. Our role shifts towards orchestrating the design process through multiple and rapid iterations, making strategic decisions while AI executes much of the technical work under our creative guidance.

This evolution demands a recalibration of the skills and sensibilities that define architectural expertise. Architects must cultivate a new form of AI literacy, enabling them to steer the generative power of these tools toward meaningful design outcomes. We become curators and strategists, guiding AI's capabilities and making critical decisions in a fast-paced environment where the production and curation of designs happen at an unprecedented speed.

The sketches and models we create in this hybrid human-AI workflow are no longer static representations but part of an engaging dialogue with AI. These inputs are more concise, layered and designed to respond dynamically to AI prompts. We produce vast amounts of data, generating and exploring a multitude of design variations, and we must develop the intuition and critical eye to identify opportunities and discard the majority that do not lead to satisfactory outcomes.





An AI-generated model alongside an AI-generated building, illustrating the transformation from abstract form to detailed architectural design.

Layers of AI-Driven Architectural Workflows: Introducing ControlNet

To navigate this new landscape, architects require tools that allow for nuanced communication with AI systems. Contemporary AI-driven architectural design transcends simple text prompts, embodying sophisticated processes that facilitate precise guidance. At the forefront of this evolution is ControlNet, a neural network architecture developed by Lvmin Zhang, Anyi Rao and Maneesh Agrawala at Stanford University. ControlNet offers a suite of trained models for specific tasks, enabling architects to guide AI using diverse layers of information in a visual form.

Integrated as part of the open-source Stable Diffusion framework, these key components form the building blocks of a hybrid AI-human architectural design lexicon, each offering unique capabilities and control mechanisms:

1. Textual prompts. The cornerstone of AI-driven design, utilising natural language with precise weighting (e.g., 'opaque:1.5' to emphasise opacity) to initiate creativity. Essential for conceptual exploration when geometric specifics are undefined, textual prompts are powerful on their own but their influence modulates when combined with subsequent workflow layers. Prompts form the basis of the generative process, not its entirety. When combined with ControlNet layers, they refine AI's interpretation of inputs, balancing designer intent with AI-generated creativity.

2. Segmentation images. Colour-coded representations using the ADE20K colour scheme – a widely adopted dataset and colour standard for scene parsing – to delineate distinct objects or regions, guiding AI in comprehending spatial relationships and semantic boundaries. Segmentation determines the 'what' through specific colour values, perimeters and locations within the image. It offers considerable latitude for textual prompts to manifest, even at high application values, and is the closest approximation to spatial prompt control. At lower values, it aids in arranging image elements without rigidly defining their perimeters.

3. Depth maps (Z-depth images). Grayscale representations of surface distances, enabling precise control over spatial relationships in AI-generated or modified images. When applied strongly, depth maps can determine geometric details with high fidelity but can also be utilised to arrange more flexible, prompt-defined geometries in space. They synergise effectively with segmentation models.

4. Edge and structure guidance. Algorithmic reduction of images to essential outlines, providing a framework for AI models to develop from overall structure to fine details and textures. Options include Canny (for highly detailed geometry), Scribble and SoftEdge (for outlining and guiding) and M-LSD (line segment detection) for straight line definition, each serving distinct architectural needs.

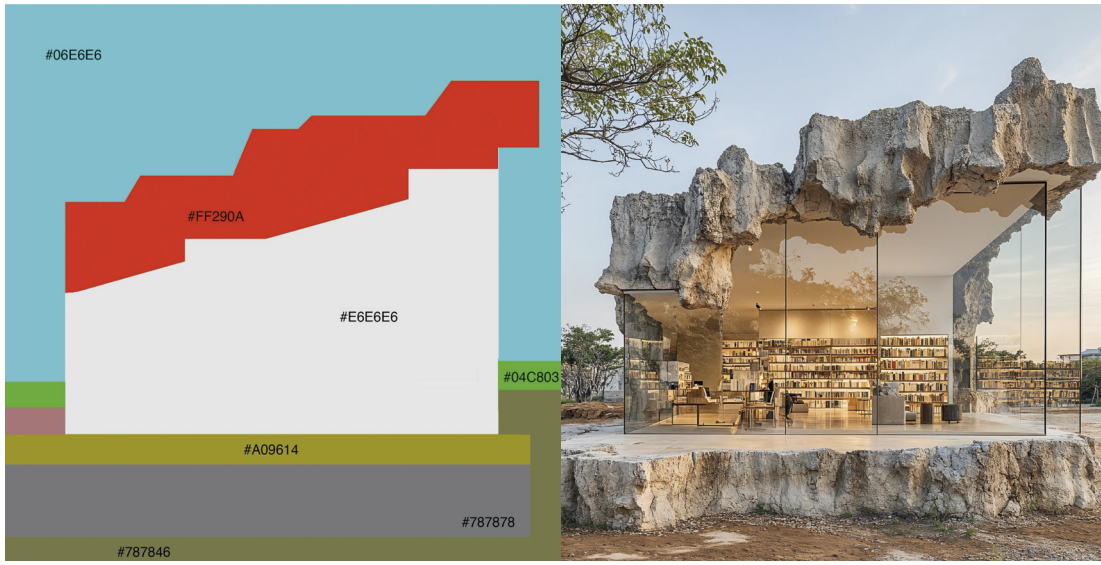
5. Image prompt adapters (IP -Adapters). A powerful method for incorporating visual references directly into the generation process. IP-Adapters integrate exemplar images, fine-tuning the model's output to align with specific architectural styles, textures or spatial configurations. This technology enables a level of control

and specificity previously unattainable with text prompts alone, allowing architects to seamlessly blend historical references with contemporary innovations in their generative workflows.

6. Inpainting and outpainting. Models enabling targeted, coherent modifications within images, crucial for exploring design variations and making precise adjustments. Inpainting allows for modification, bridging geometries or filling areas, while outpainting extends images, with significant implications for connection and articulation in design resolution and the addition of contextual information.

Collectively, these components forge a new architectural glossary within AI-assisted design, providing architects with diverse methodologies to effectively communicate and collaborate with AI systems. The integration of this hybrid language opens new avenues for creative exploration, design expression and precision in the architectural field.

As we harness these sophisticated tools, we must consider: How does this new visual language reshape our design thinking? What new possibilities emerge when we can communicate architectural intent through data layers as fluently as through sketches? And how might this level of control over AI-generated designs influence the future of architectural representation and communication?



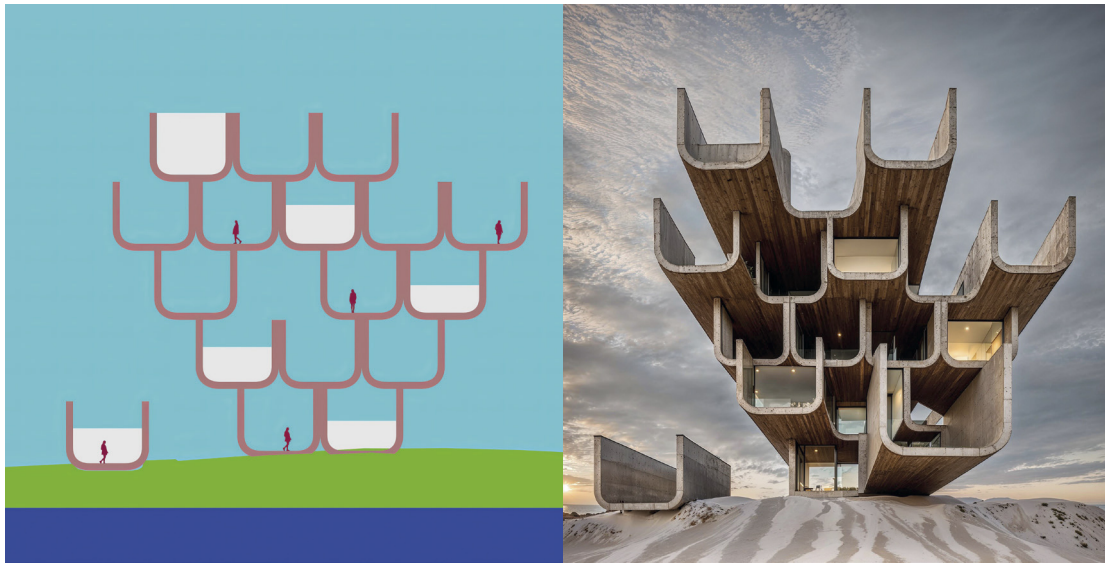
Layers of ControlNet showcasing segmentation input for precise design generation.

Human-AI-Driven Workflows: Case Studies

The integration of advanced AI techniques into architectural practice has given rise to innovative workflows that blend human creativity with machine intelligence. These case studies illustrate how the AI-driven processes we've explored are reshaping architectural design, offering new avenues for exploration, refinement and communication of ideas.

1. Conceptual segmentation-based 3D modeling
Workflow: AI segmentation → Rhino 3D modelling → Human-scale integration

This approach leverages AI segmentation using the ADE20K colour scheme to bridge 2D concepts and 3D architectural modelling. Architects create conceptual models in Rhino 3D where the layers' colours precisely match the segmentation scheme, guiding both geometric boundaries and semantic information. This method allows for the rapid translation of AI-generated ideas into spatially aware 3D models. The integration of human figures helps maintain a clear sense of scale, while the use of the 3D modeler's 100% 'Self-Illumination' feature preserves the clarity of segmented areas. For enhanced guidance, this workflow can be combined with a Z-depth map directly generated from the 3D model, further refining spatial relationships.



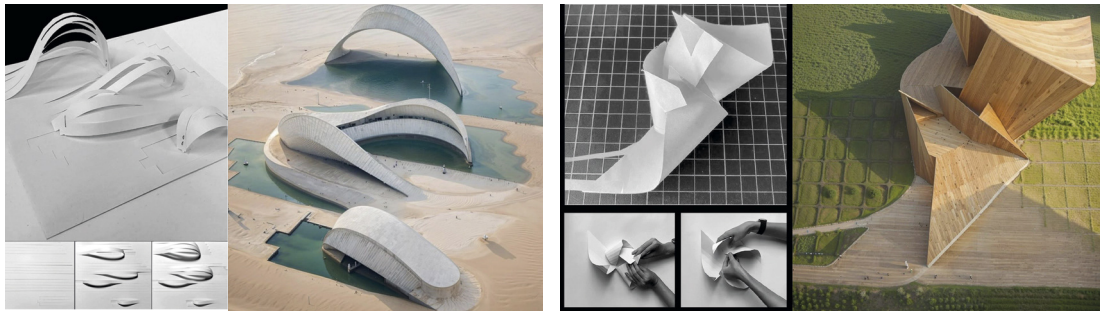
(Left) Segmentation image from Rhino3D.
(Right) AI-generated architectural design informed by segmentation.



(Left) AI-generated contextual image. (Middle and right) Architectural volumes integrated into the existing environment using ControlNet.

2. Seamless design integration through inpainting
Workflow: Precision masking → AI inpainting for transitions and connections

In this workflow, the crucial first step is to set the mask with extreme precision, as the inpainting model considers pixels adjacent to the boundary when generating new content. This high level of accuracy in masking ensures that the AI can effectively 'read' the surrounding context and extend it seamlessly. This technique is particularly valuable in renovation and urban infill projects, allowing for rapid exploration of integration strategies while maintaining design coherence. By carefully defining the areas for inpainting, architects can guide the AI to produce contextually appropriate transitions, resulting in designs that appear organic and well-integrated with their surroundings.



(Left) Paper model.
(Right) AI-generated visualisation using depth mapping inferred from the model.

3. Conceptual physical model to spatial representation
Workflow: Physical model photograph → AI depth mapping → Guided Visualisation

This workflow bridges the tactile world of physical modelling with advanced AI visualisation techniques. Architects begin by photographing their physical models, which are then processed using AI depth estimation models. These tools infer the depth information from the 2D photograph, creating a Z-depth map. This depth map is then used to guide AI-generated prompts, transforming the simple model photo into a sophisticated architectural visualisation. The inclusion of human figures in the model's picture – these can be handmade – helps establish scale and brings life to the visualisation. This process allows architects to rapidly iterate between physical modelling and digital representation, combining the intuitive, hands-on approach of model-making with the power of AI-driven visualisation.



Existing textures integrated via IP-Adapter.



New design materiality using AI inpainting and dual IP-Adapters.

4. Differentiated localised materiality using advanced IP-Adapters
Workflow: Inpainting + dual IP-Adapters for context and new design

This sophisticated approach combines inpainting with dual IP-Adapters to address material integration in any given context. One adapter focuses on existing textures and contextual characteristics, while the other on the new design's aesthetic. This allows for rapid exploration of materiality options that harmonise with the surrounding environment while maintaining design integrity. This workflow enables architects to create designs that are both innovative and contextually sensitive, facilitating a more nuanced approach to urban integration.

These case studies demonstrate the practical application of AI-driven workflows in enhancing architectural design processes. By synthesising various AI techniques with traditional design skills, architects can explore design possibilities more efficiently, communicate ideas more effectively, and bridge the gap between conceptual thinking and practical implementation. As these technologies continue to evolve, they promise to play an increasingly integral role in shaping the future of architectural practice.

The Architect as the New AI Artisan

As AI and no-code products become increasingly accessible, the true artistry in architectural design is shifting towards the creation and manipulation of workflows. The workflow itself becomes our craft, allowing us to create sophisticated design processes that blend human creativity with AI capabilities. In this new paradigm, architects are not just designers of buildings but artisans of novel digital workflows and curators of AI models. This shift requires architects to develop a hybrid language, where intuition, creativity and technical knowledge combine seamlessly with computational power. Moreover, the training of specialised AI models, particularly through techniques like low-rank adaptation (LoRA), emerges as a new frontier for architectural creativity and differentiation. By becoming authors of our own models, we regain control over our design tools. The creation of these personalised AI models involves curating image sets, developing thoughtful tagging systems and fine-tuning parameters – essentially crafting our own 'secret sauce'. This process becomes an extension of the architect's creative expression, allowing for the infusion of personal style, regional influences or specific design philosophies into the AI tool itself.

By developing specialised models addressing structural considerations, contextual sensitivities or unique aesthetic visions, architects can create methods that embody their expertise and design philosophy, standing out in an increasingly AI-driven field. The power of this approach lies in its specificity, contrasting with the averaged outputs of large, general models and offering a competitive edge in both creativity and efficiency. Architects might develop suites of models, each fine-tuned for specific aspects of the design process – from conceptualisation to sustainability analysis.

Looking ahead, the ability to develop and wield these personalised AI models may well become a defining skill for architects of the future, marking a new era where computational expertise and creative vision are inextricably linked.

Architecting Tomorrow: The Dialogue Evolves

The integration of AI-driven workflows will fundamentally reshape architectural practice, with implications that extend far beyond the realms of conceptualisation and visualisation. ControlNet's visual data layers in combination with diffusion models have emerged as essential tools, enabling architects to communicate with machines through images as naturally as they do with words. In a profession deeply rooted in visual expression, leveraging data in this manner represents an unprecedented shift. However, the true potential of these AI models will be amplified when merged with structural analysis, efficiency simulations, building information modelling (BIM) and other parametric tools, creating increasingly complex workflows that span all stages of a project. This integration promises to revolutionise not just ideation and concept development, but also advanced design stages and detailing and manufacturing processes.

As these workflows grow in complexity, we can anticipate the development of user-friendly interfaces that may oversimplify the underlying processes. While this accessibility is beneficial, it also underscores the critical need for architects to remain in charge of these workflows, understanding their intricacies and potential pitfalls. Our role evolves to become not just creators of buildings, but architects of these sophisticated, multi-layered processes.



Concept design integrating a personalised LoRA model in the workflow together with ControlNet segmentation and depth layers.

However, as with any technological advancement, there is a dichotomy in its adoption and impact. While progressive architectural practices will harness these new tools to innovate and distinguish themselves, others may resort to automating existing models, replicating familiar designs with greater efficiency but less originality. This could represent the majority, as many stakeholders may not demand the specificity and innovation that advanced workflows enable. The risk is a proliferation of homogenised architecture, where the potential of AI is used merely to expedite the replication of the status quo. Yet, for those willing to embrace the full spectrum of AI's capabilities, there lies an unprecedented opportunity to redefine the architectural landscape by creating responsive designs imbued with a distinct identity.

Engaging in this critical dialogue is essential as we redefine our paradigm, setting new standards that liberate us from past and sometimes restrictive practices. By utilising increasingly natural language interfaces and AI-driven tools in the design process, we are 'talking' in the most literal sense – engaging in a continuous conversation that reshapes our profession. This new discourse becomes a collaborative partnership between architects and AI, co-designing in synergy. Paradoxically, by adopting AI-driven methods, we deepen our connection to the human element and our foundational principles by focusing on what truly matters.

The conversation between architecture and AI is just beginning, and it promises to be one of the most transformative shifts in our profession's history. Architecture at 60fps is more than a metaphor for speed – it represents the vast range of possibilities we can explore instantly. It highlights the density of information we process and the rapid pace at which we can generate, curate and iterate designs. Architecture, from now on, must embrace this evolving dialogue, one that will redefine the very fabric of our built environment.



Sculpting with Inference

ChatGPT-4 + Processing 4.0 + ComfyUI

Carlos Bañón

DOI: 10.4324/9781003726500-13

JOURNEY

In contemporary architectural practice, the integration of computational tools and AI is increasingly redefining the design process. Architects are no longer limited to traditional methods; instead, they can now integrate the power of both procedural design and generative AI to explore new creative territories. This novel workflow exemplifies such an approach, where the intersection of procedural and generative design is explored to achieve both precision and creative freedom.

The workflow begins by developing a procedural base using Processing 4.0 code to generate schematic geometries. These geometries serve as the foundation for the subsequent AI-generated images. This procedural method allows architects to establish initial parameters - geometric system, scale and composition - while maintaining flexibility for AI-driven exploration.

Semantic segmentation plays a pivotal role by establishing boundaries that guide the design process. These boundaries create a clear framework within which the AI operates, facilitating controlled and intentional image generation. By directly drawing with segmentation colours, architects can create geometric boundaries imbued with meaning and context, setting the stage for further refinement through AI.

The collaboration between human and AI in this procedural design and generative AI creates a distinct and unique balance. While the procedural approach emphasises structure and order, the generative side introduces variation and complexity. AI manages the articulation, detailing and resolution of elements, while human input dictates the overall scheme, materiality and atmosphere through written prompts. This synergy extends the boundaries of conventional architectural generation, fostering continuous exploration and innovation within a well-defined framework, ultimately leading to deliberate and inventive architectural expressions.

WORKFLOW

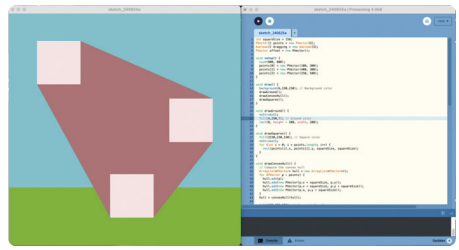
STEP 1

We begin by creating a system that defines our architectural design. The focus will be on generating a façade composed of three squares that are mathematically wrapped by a polygonal shape (convex hull) defined by the squares' corners. To achieve this, we will use Processing 4.0 as our procedural tool, allowing us to easily generate variations of our geometries. To create the code in Processing 4.0, we use a prompt in ChatGPT-4: Generate an 800x800px Processing 4.0 sketch showing three

draggable 150x150px squares. Use all the square vertices to create a convex hull polygon. The ground is a full-width, 200-pixel-high rectangle at the bottom. Draw order: background (RGB 6,230,230), the ground (4,250,7), the convex hull (180,120,120), and the three squares (230,230,230). Figures have no stroke and no transparency.

It is possible that the script has unexpected bugs. If so, copy the error from the console and ask ChatGPT to fix it using the prompt: 'Fix this error: error message. Provide full code.'

Alternatively, the base image can be created in Photoshop or other design software using the same colours indicated above.

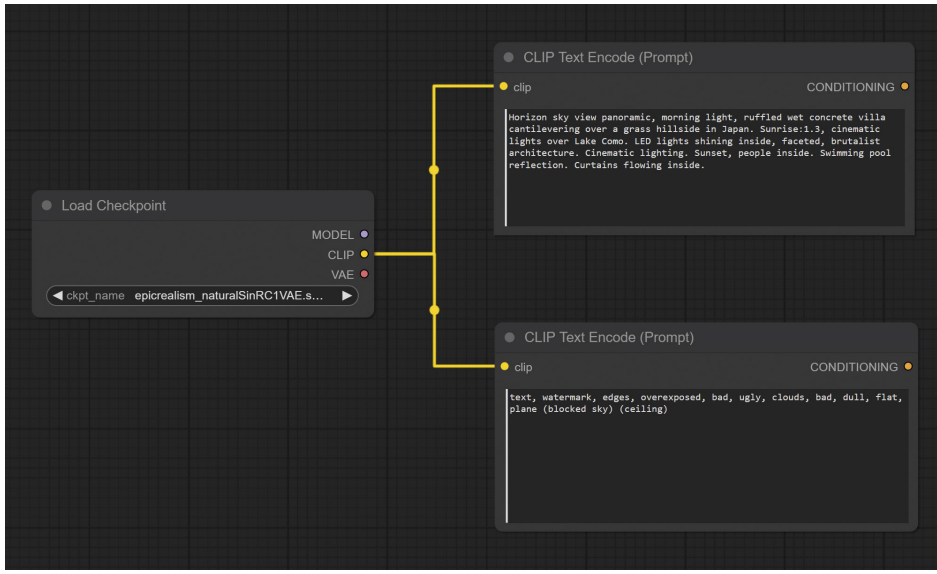


Creation of a basic geometric system in Processing 4.0 with the help of ChatGPT.

STEP 2

With the procedural tool established, we turn our attention to the generative aspect, utilising ComfyUI to run Stable Diffusion. For the first set of nodes, we need to establish the diffusion model that will be used to generate the image. We will create a 'Load Checkpoint' node and use 'EpicRealism.safetensors' (available at *https://civitai.com/models/25694/epicrealism*), as it is highly capable of creating realistic architectural images.

Next, we create two containers for positive and negative prompts by double-clicking on the canvas and selecting ‘CLIP Text Encode’ for both. For the positive prompt, we will leave it customisable depending on the desired architectural style. For the negative prompt, we will use typical negative prompts such as ‘text, watermark, edges, overexposed, bad, ugly, clouds, bad, dull, flat, plane (blocked sky) (ceiling)’.



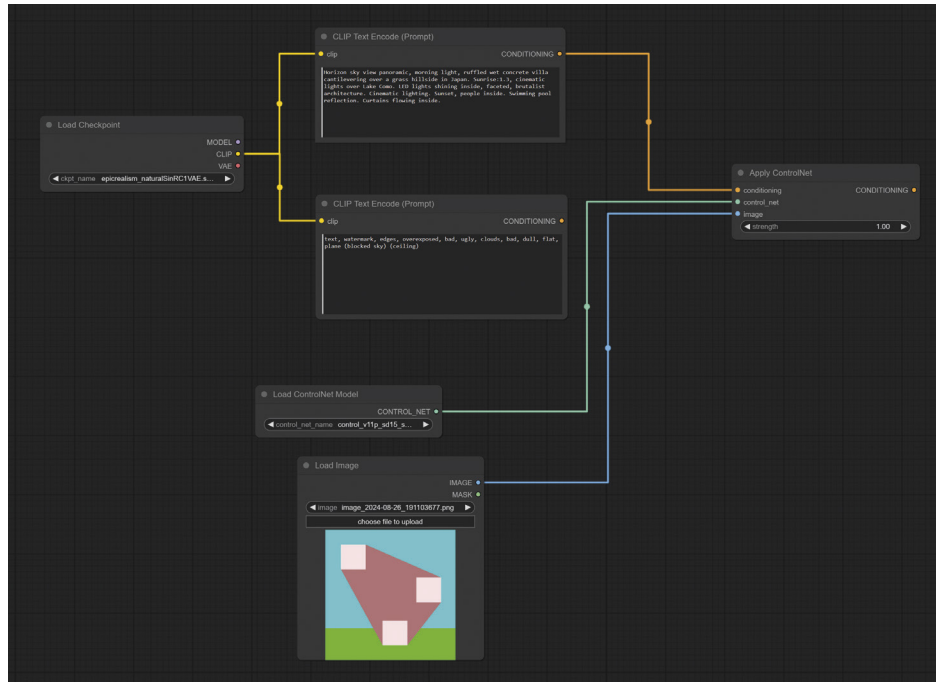
Selecting the checkpoint model and prompts.

STEP 3

Once our positive and negative prompts are set, the next step is to integrate ControlNet into the workflow to provide more precise control over the image structure. We begin by loading the ControlNet segmentation model through the ‘Load ControlNet Model’ node. In this case, we use the model ControlNet-v1-1-segfp16 (available at *https://huggingface.co/comfyanonymous/ControlNet-v1-1_fp16_safetensors/tree/main*), which is effective for guiding the diffusion process based on segmentation or layout images. This model ensures that the generated architecture adheres closely to the intended structure and composition.

After the ControlNet model is loaded, we introduce an image node via the ‘Load Image’ node. This node will take a guiding image to be used as a segmentation map, which will outline the main elements of the scene, such as the villa and surrounding landscape. We will use the image from our procedural tool in Processing 4.0 using the snapshot feature of your computer.

Finally, everything converges in the ‘Apply ControlNet’ node, where the positive prompt, ControlNet model, and guiding image come together. We set the strength to 1.0, ensuring that the diffusion process is fully guided by the ControlNet model and the input image.



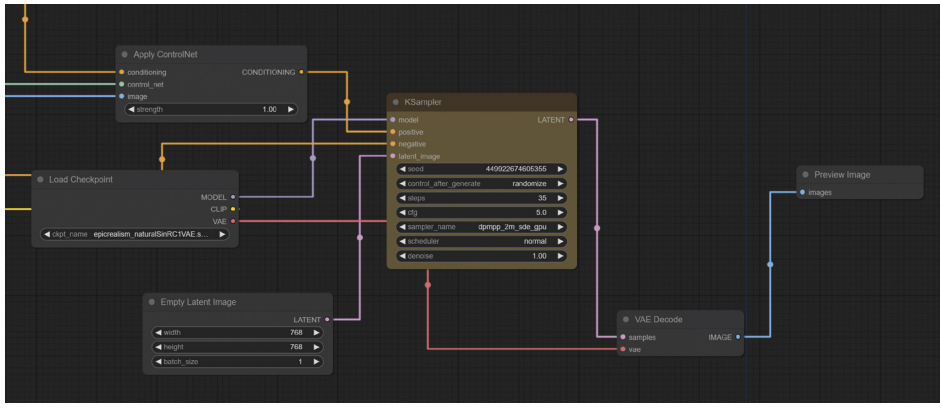
Applying ControlNet to guide the diffusion model with a segmentation map.

STEP 4

Once ControlNet is integrated, the next step is to set up the image generation process using the KSampler node. The KSampler is crucial for defining the diffusion process parameters and converting the conditioning data into a latent image. We begin by connecting the 'Apply ControlNet' node to the KSampler's conditioning input. The KSampler will use the 'dpmpp_2m_sde_gpu' sampler with a denoise strength of 1.0. The classifier-free guidance (CFG) scale is set to 5.0, providing a balanced influence of both the prompts and the model's learned features. The number of steps is configured at 35 to ensure a detailed output without excessive computation time, but it can be reduced to 20 without significantly compromising the quality.

In parallel, we prepare an empty latent space using the 'Empty Latent Image' node, set to a resolution of 768x768px with a batch size of 1. This latent space serves as the canvas where the diffusion model will generate the image. Connect the latent space to the KSampler's latent image input to initiate the generation process.

After sampling, the output is sent to the 'VAE Decode' node, which converts the latent image back into a visible image format. Remember to connect the VAE component from the 'Load Checkpoint' initial component. Finally, the decoded image is previewed using the 'Preview Image' node, allowing us to see the generated architecture based on our prompts, ControlNet guidance and diffusion settings. Now, click on 'Queue Prompt' to generate your first image.



Setting up the KSampler.

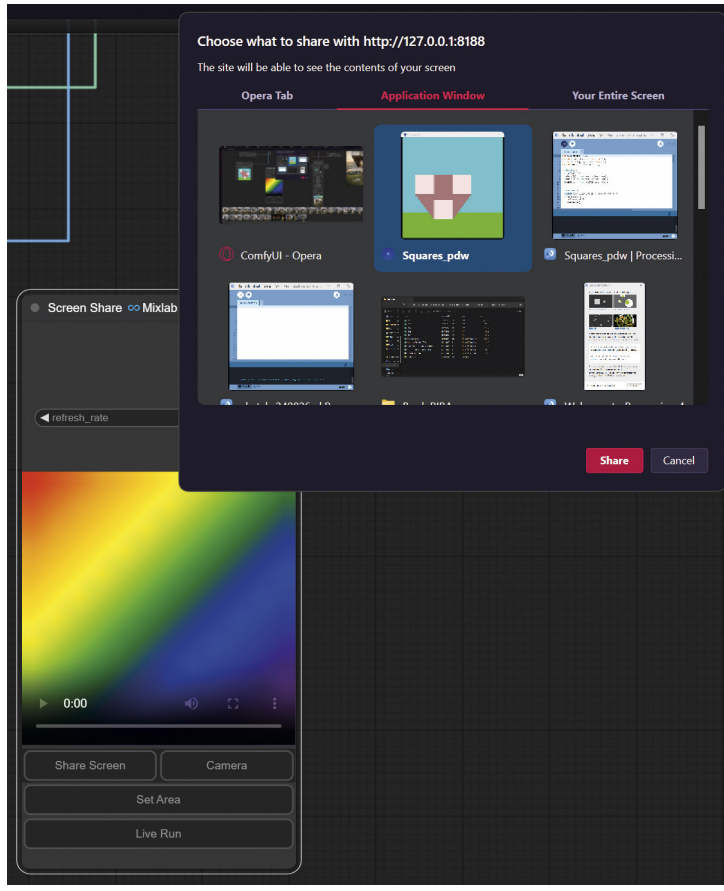


Generating our first image using segmentation.

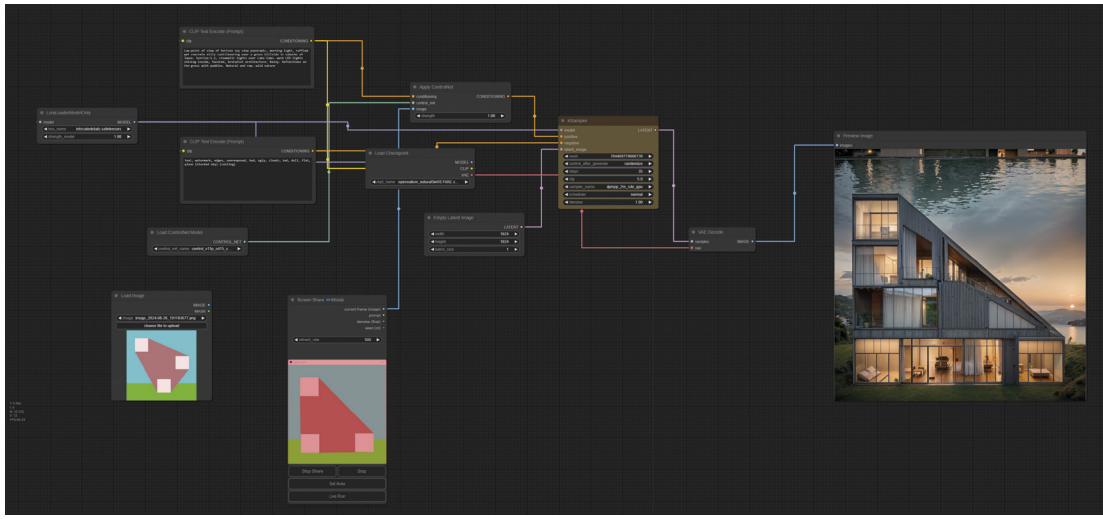
STEP 5

For the final step, we enable real-time data transfer from Processing 4.0 to Stable Diffusion using the MixLab 'Screen Share' node (available at *https://github.com/shadowcz007/comfyui-mixlab-nodes*). To set this up, click on 'Share Screen' and select 'Application Windows', where you'll see your Processing 4.0 interface. Once selected, the live image from Processing 4.0 will be seamlessly transferred to ComfyUI. Activate 'Live Run' to initiate the data flow. If you want to exclude the top bar of the captured image, you can use the MixLab node's area selection feature to focus only on the desired section. The image output from the MixLab node will then connect directly to the ControlNet component, replacing the static image used earlier in the workflow.

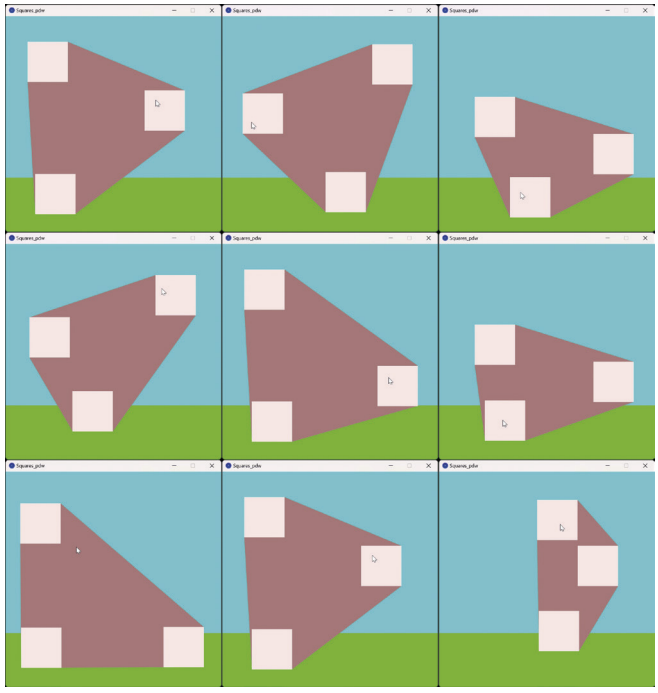
Finally, replace the 'Preview Image' node with the 'Save Image' node at the end of the workflow to automatically store the generated image locally on your hard drive.



Using the 'Screen Share' feature from MixLab nodes.



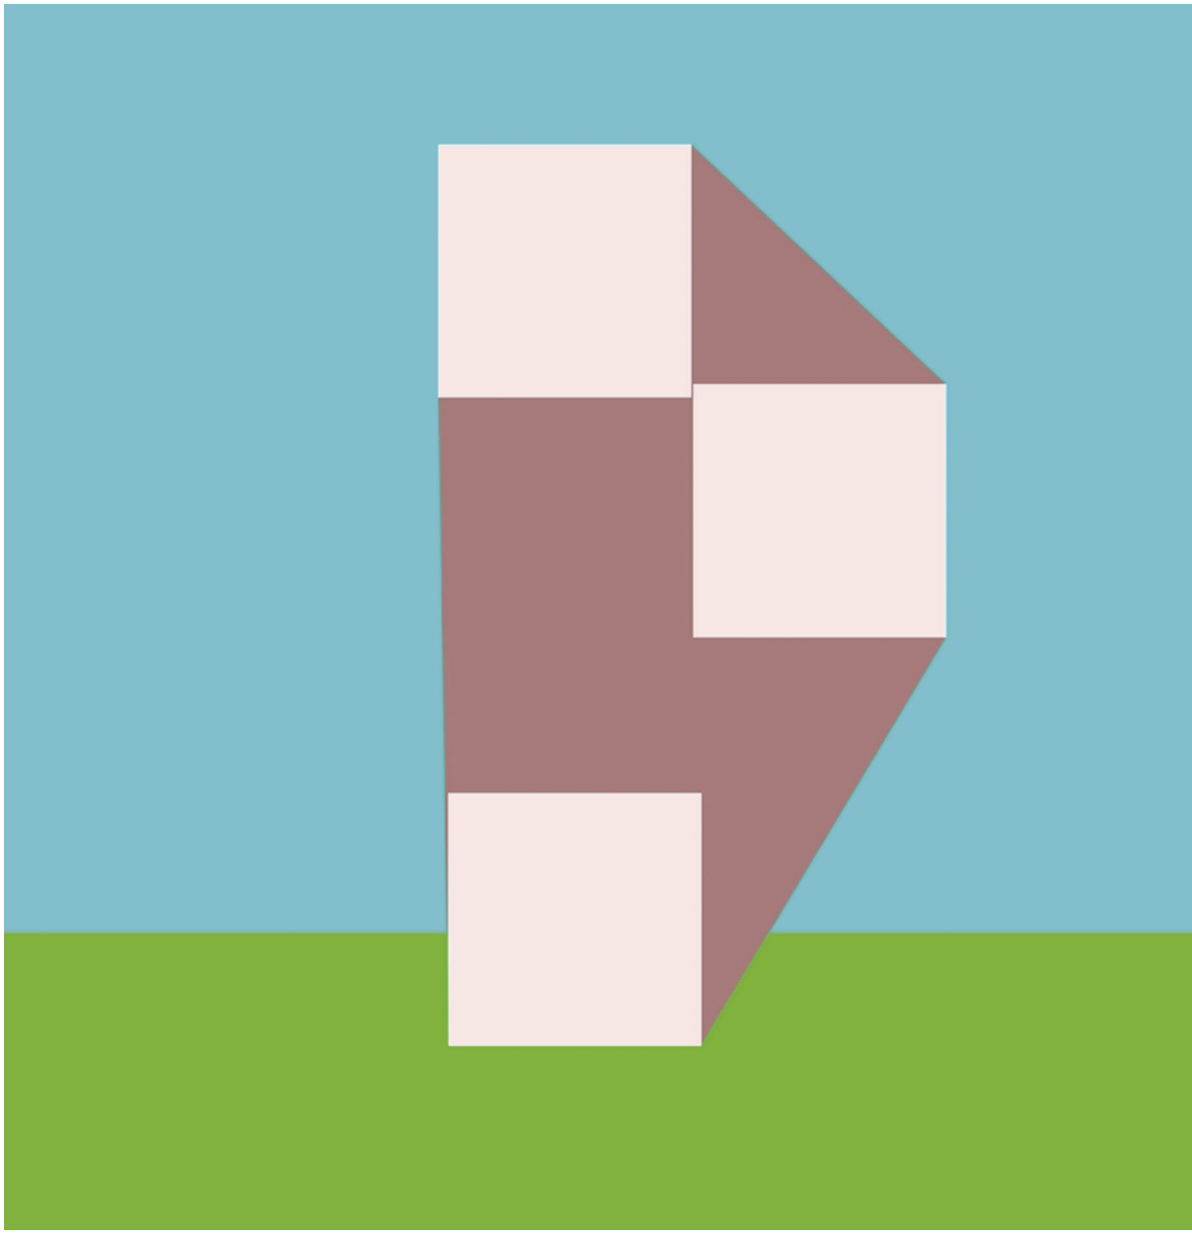
Full workflow.



Grid of nine procedural images in Processing 4.0.



Grid of nine generated images using the full workflow.



Final result - from procedural geometry to AI-generated realistic image.



TIPS

Have a vision, then embrace iterative exploration	Devote time to systematic experimentation with AI outputs. Allow each iteration to challenge conventional design constraints while remaining firmly rooted in a clear, overarching architectural vision.
Integrate timeless principles	Combine enduring architectural tenets with contemporary computational techniques. Let your mastery of structural analysis and digital fabrication guide AI outputs, fostering a seamless dialogue between classical design principles and modern technology.
Train your eye, cultivate a swift curatorial process	As vast quantities of data are produced rapidly, develop an efficient system for evaluating iterations. This disciplined approach enables you to swiftly identify innovations that truly enhance your overall vision.

What does the future role of the architect or designer look like?

I envision the future role of architects and designers in a transformative synthesis where AI seamlessly integrates with advanced tools, enhancing efficiency, innovation and responsiveness to environmental and cultural contexts. While AI models will surpass our computational capacities, our role remains that of curators and visionaries. Strong design fundamentals, clear architectural vision and experiential understanding will allow us to harness AI's power while preserving the human element. This balance of machine precision and human intuition will redefine the built environment as a dynamic, integrated narrative, ensuring that creativity and sensibility remain central to architectural innovation.



Real-Time AI-Enhanced Digital Collaging-to-3D

Krita + Stable Diffusion LCM Plugin + TripoSR v2.5 + Rhino

Hamza Shaikh

DOI: 10.4324/9781003726500-14

JOURNEY

In my pursuit to preserve traditional artistic methods while gaining enhanced AI-powered control over design composition, I demonstrate a real-time AI inferencing tool that combines procedural prompting with digital collaging. This approach transforms ideation into a dynamic visualisation exercise rather than a static workflow. Using two open-source tools - Krita and a plugin from *interstice.cloud/plugin* - the workflow integrates Stable Diffusion's image-to-image and ControlNet functions seamlessly within Krita's digital art environment.

Designing, visualising and iterating directly within a chosen key view, without modelling detailed 3D geometry, is a game-changer for early concept stages. Digital collaging enables rapid exploration of ideas through the compositing of materiality, form and precedent from diverse references. However, traditional collages can feel crude, lacking spatial depth, atmosphere and immersion.

By introducing AI, rough collages are dynamically refined through real-time prompts. As elements are moved, rotated or distorted, the plugin abstracts and integrates them based on the prompt, allowing users to control how much the collage shifts from literal representation to abstract synthesis. Key adjustments are made via four simple controls: strength, prompt, negative prompt and seed.

Advanced users can load custom models and fine-tuned low-rank adaptation (LoRA) models, but even default settings provide a seamless, intuitive experience, aligning the AI's output with creative intent.

The final step will show how to transform a single image of the completed visual into a 3D model using Tripo3D's image-to-3D model. This model is then imported into Rhino to be interrogated and potentially optimised with additional scripts.

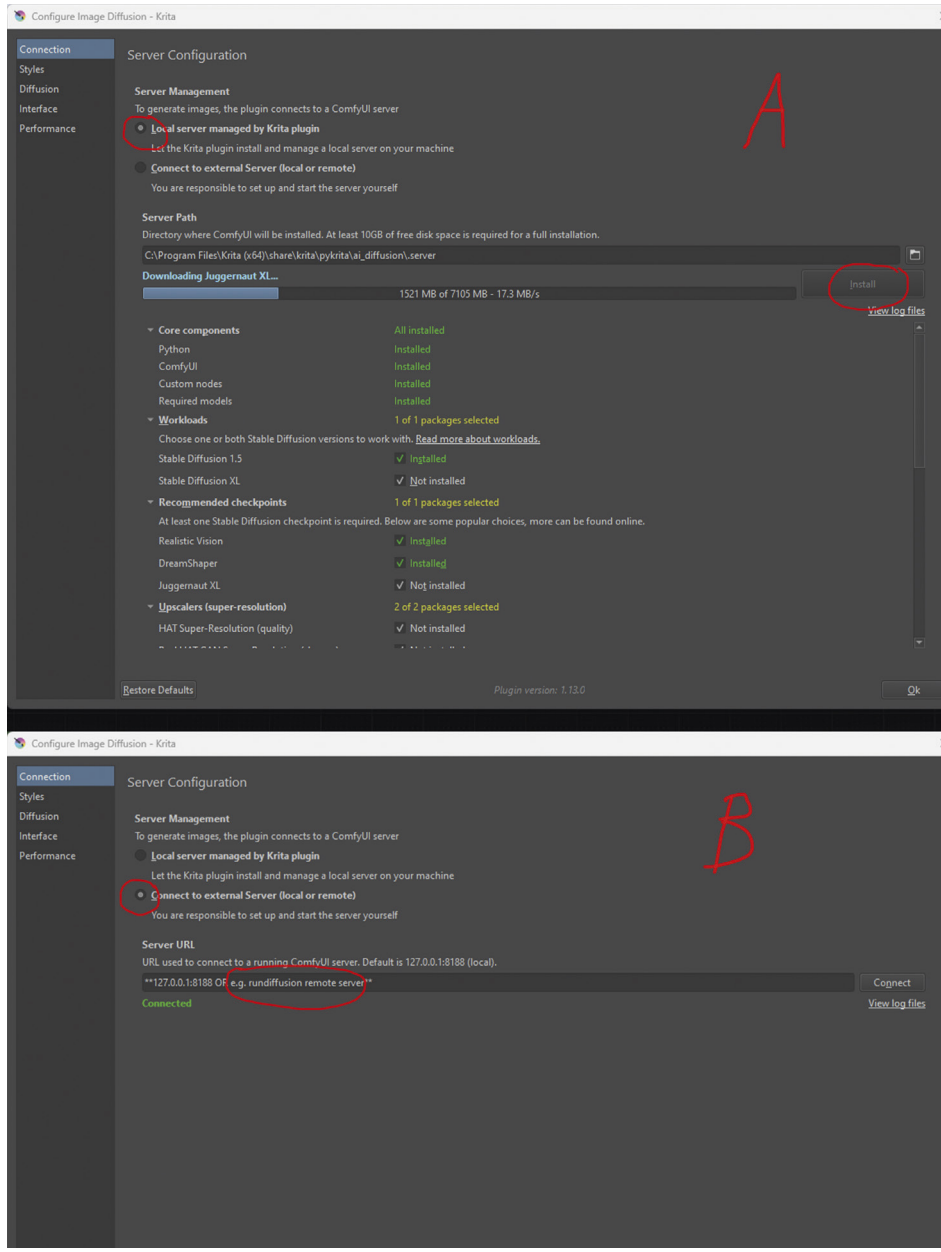
WORKFLOW

STEP 1

To begin you need to install the Krita drawing software to your computer and then the AI Diffusion plugin to configure the tools together. Keep in mind before downloading that you will have two options after downloading and installing tools: A. To run the AI program locally (if you have at least 8GB of video random-access memory (VRAM) and a higher spec graphics card e.g. NVIDIA RTX 4000 series) or B. To run it using an external remote server connection if you do not have a high spec graphics processing unit (GPU).

In the case of option A, ensure you have enough memory to store the large amount of model and file data also. In the case of option B, Run Diffusion offer a simple application programming interface (API) connection to run the program simply through an internet connection.

More details for downloading and installing steps can be found here:
Krita and AI Diffusion install: *https://www.interstice.cloud/plugin* -
Run Diffusion remote server API connection: *https://learn.rundiffusion.com/* -
krita-plugin/ -



Install and configure Krita and AI Diffusion plugin establishing server.

STEP 2

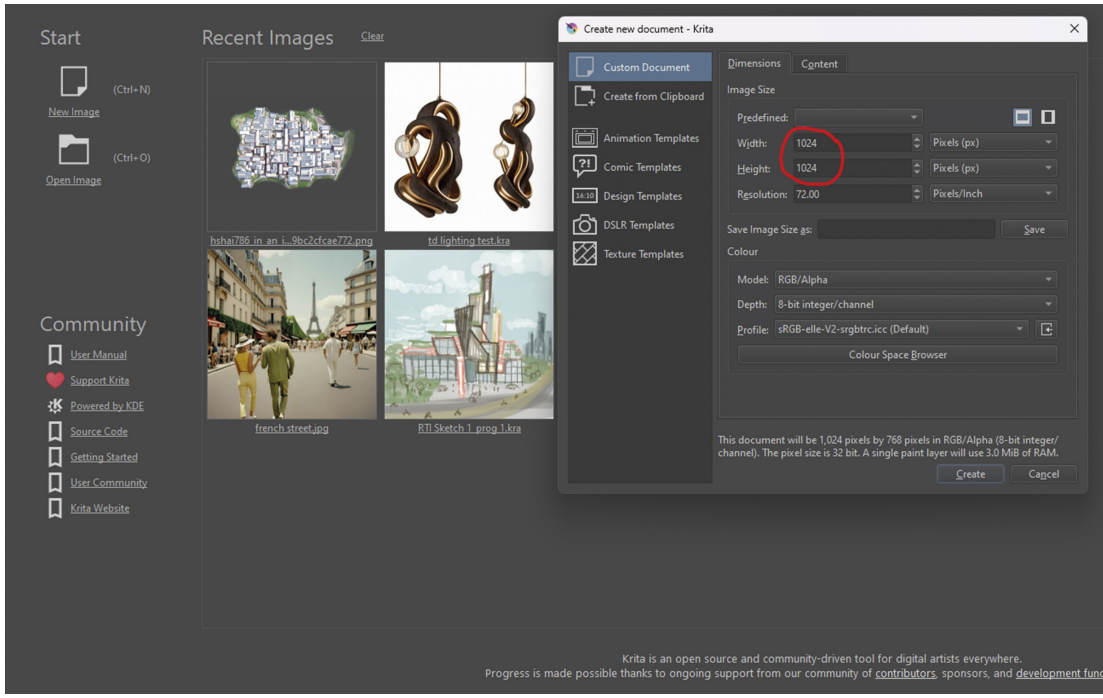
Before starting an image generation exercise, it's essential to understand the 'golden dimensions' 512, 768 and 1024 – key pixel values used in training diffusion models. For optimal results, use at least one or combine two of these dimensions.

Examples:

- Square: 1024x1024px for high-quality, balanced outputs.
- Landscape: 1024x768px for rectangular compositions.
- Fastest response: 512x512px to minimise GPU load.

Your choice depends on your goals and the model. XL models typically require at least 1024x1024px. Variations like 1366x1024px are valid if they include a golden dimension, ensuring high-quality results.

In setting up my Krita canvas to begin, I will use 1024x1024px for a 1:1 ratio, which suits my needs and supports XL models.

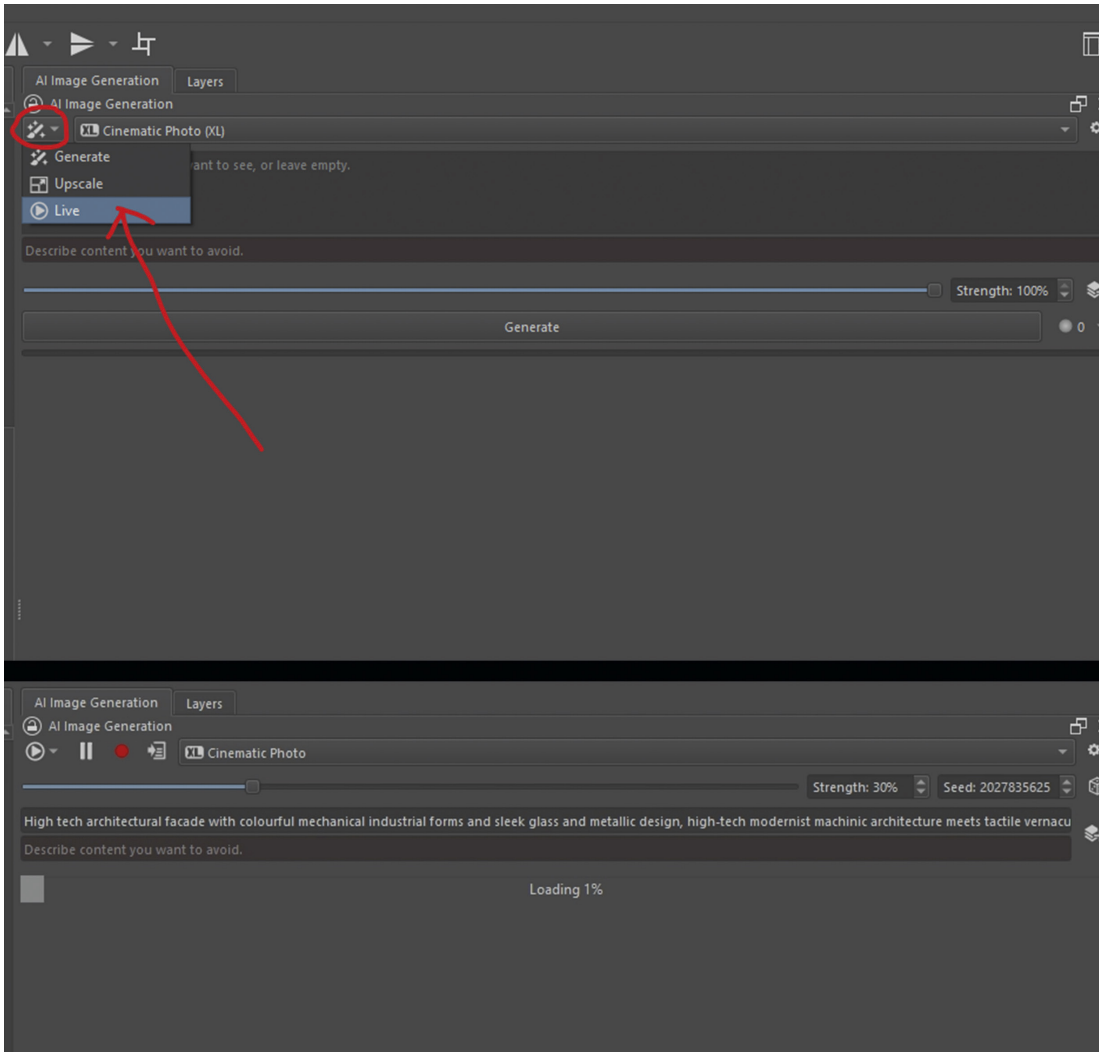


Set width and height as needed using 'golden dimensions'.

STEP 3

Assuming you have installed and configured the AI Diffusion plugin into Krita with instructions and links from Step 1, you can now enable ‘Live’ mode from your configured ‘AI Image Generation’ docker tab.

After enabling ‘Live’ mode, begin writing your prompt and prepare your collage ‘pieces’ or begin searching for your images.

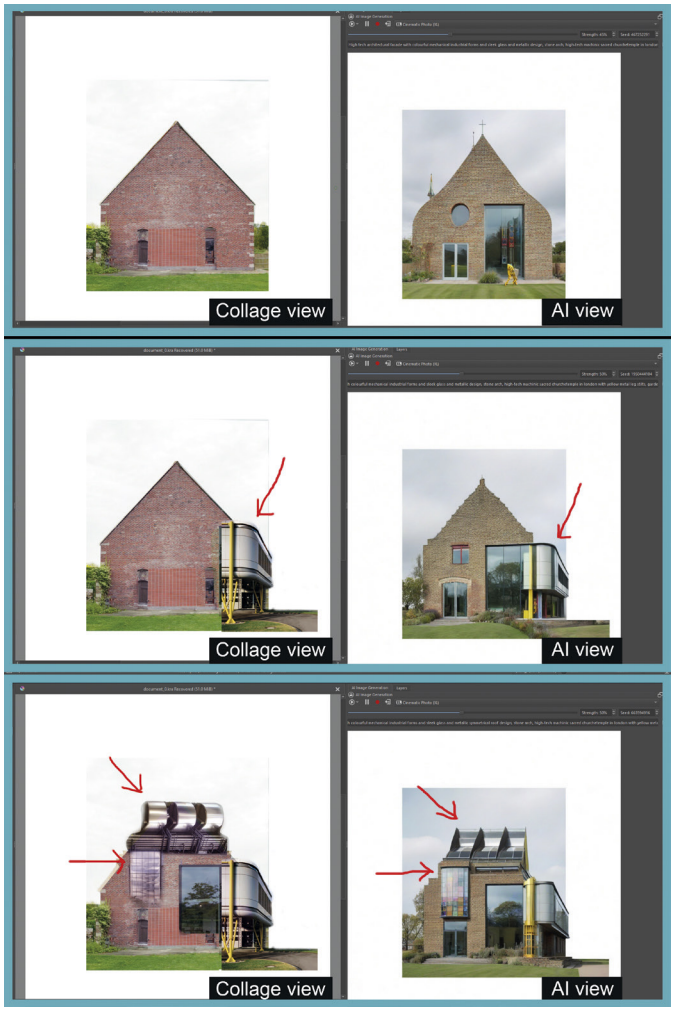


Enable ‘Live’ mode and write the first pass of your prompt and collect your images.

STEP 4

Start your digital collage with ‘Live’ mode enabled. This allows your prompt to actively interpret the digital canvas on the left, translating it into a continuous image generation feed displayed on the right. Any new collage pieces you add are instantly processed in real time by the AI.

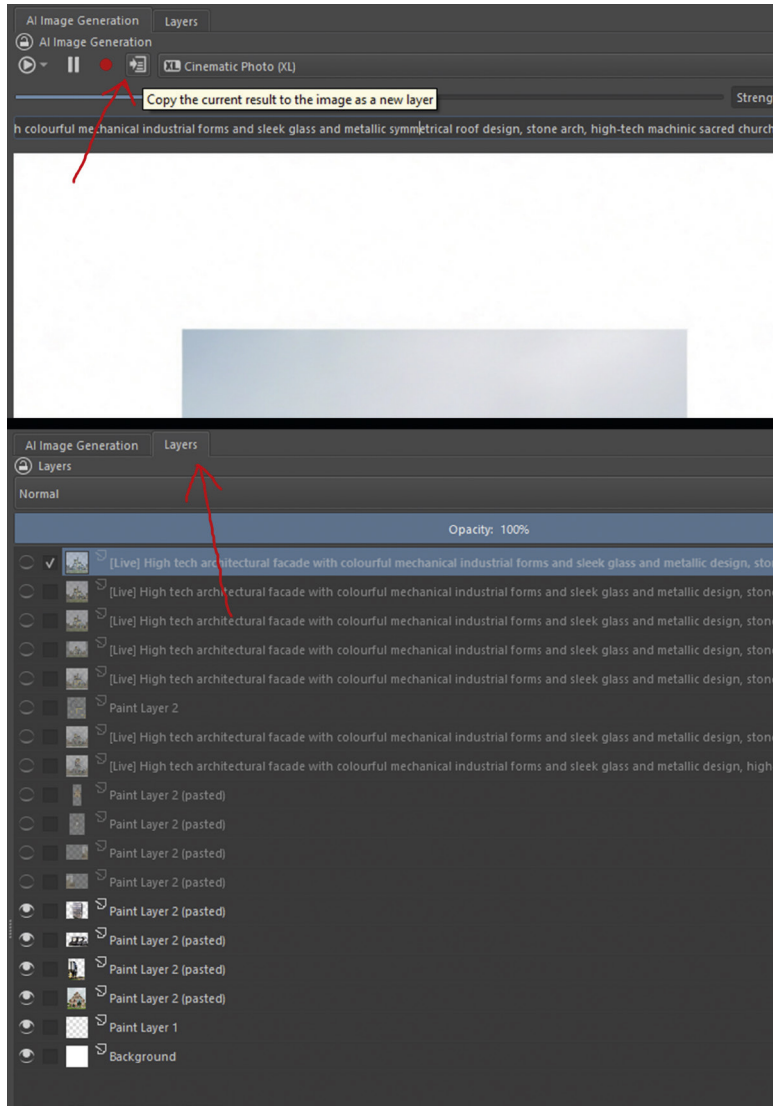
Refine your prompt as you work to align it with your design vision. Adjust the ‘Strength’ parameter in the top-right corner to control how closely the AI adheres to your collage. For unique variations, click the ‘Dice’ icon to randomise the seed value, generating a completely new version of the image.



Begin collaging on your Krita canvas, adjusting prompt, strength and seed.

STEP 5

When you generate an outcome you like, click the ‘Copy’ icon in the top-left of the docker tab to automatically save the image to a new layer. This creates an iterative record of your design process, enabling you to revisit or refine earlier versions. Think of it as a ‘Bake’ command in Rhino, letting you build upon chosen outcomes seamlessly.



Copy outcomes to new layers to shortlist and iterate.

STEP 6

Enhance your collage by using digital art tools for precise control, such as drawing more design details, painting the sky or adding .pngs of people to bring the scene to life.

Ensure your prompt is updated to account for any new elements you introduce, so they integrate seamlessly into the composition.



Draw on top and collage more elements like characters.

STEP 7

Experiment with strength, seed, prompts and collage details to explore your design fully. Higher strength values encourage design flexibility, while lower values offer greater control for refinement. Pro tip: save effective seed values with notes in a separate document for easy reuse.



Explore options and expand your design thinking.

STEP 8

Once you've finalised your design or options, upscale and post-edit the result in Photoshop or your preferred software to address any issues and add a final polish. This step highlights the evolution of your design process – staying true to the initial collage while expanding and refining it faster and more fluidly than would be possible without AI.







Final touches - post-edits.

STEP 9

Using Tripo3D accessed from *tripo.ai/app* you can upload your image to convert it into a 3D mesh model. Try to use an image with a blank background so that the subject for 3D conversion is clearly visible. Additional details like shrubs or people should be removed before conversion using Photoshop, for example.



Image-to-3D conversion with Tripo3D.

STEP 10

You can choose from multiple 3D file formats for Tripo3D to output. I chose an .obj file format and imported this into Rhino. Once in Rhino, you can see the imported model is a closed mesh and can be further developed. I use a ‘Clipping Plane’ command to reveal the detail and extents of the predictive 3D generation by the Tripo3D model.

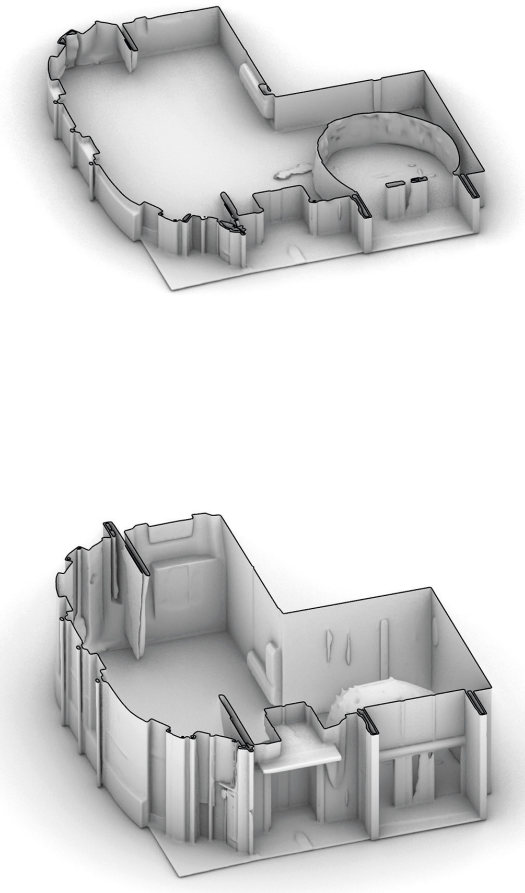
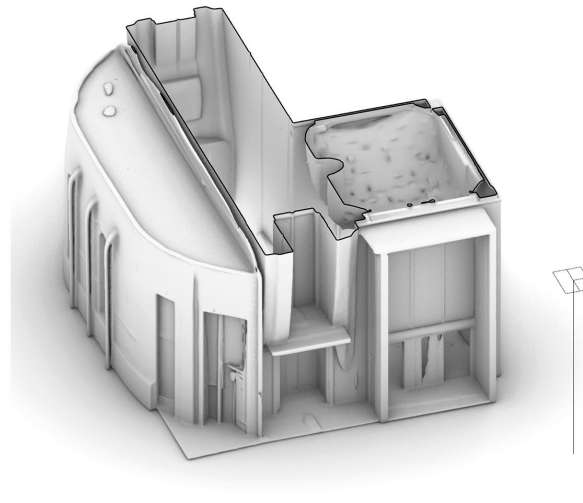
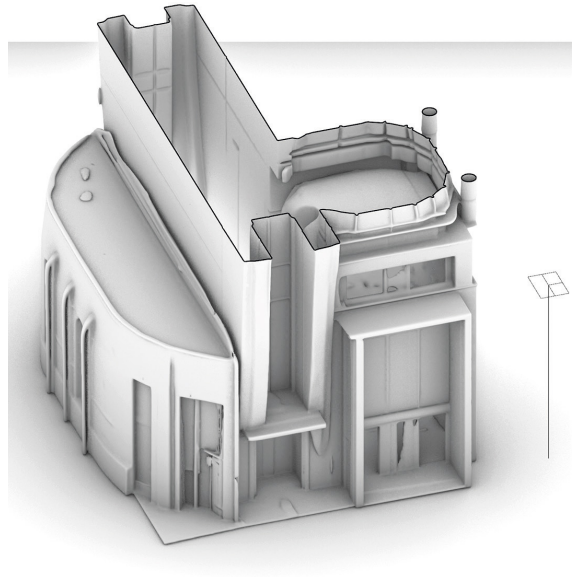


Image-to-3D conversion with Tripo3D.



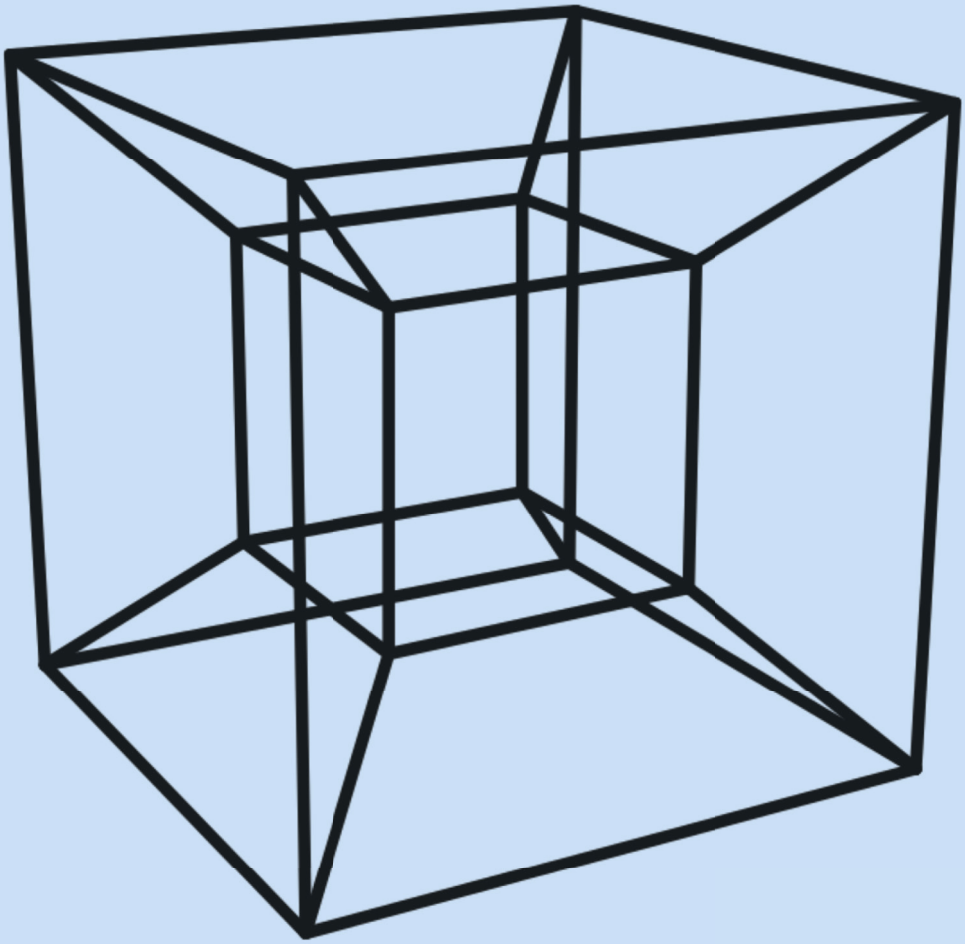


TIPS

Start simple	Keep your initial prompt minimal, then gradually refine it - adding details step-by-step. This iterative approach helps you see exactly how your words shape AI-generated images, strengthening your mind-machine connection and giving you more control.
Never lock yourself into a single AI tool	They're always evolving, so experiment with multiple open-source tools and mix them into your workflows for maximum flexibility.
Maximise your creative fingerprint	Blend AI with your own unique input. Take a hybrid approach by integrating your own drawings, original works or carefully crafted prompts in the AI generation process. If you're training a model, use your own data and fine-tune it with manual tagging to maintain creative control.

What does the future role of the architect or designer look like?

As AI-assisted coding and scripting become increasingly accessible, the horizon for technological innovation appears boundless. In a world where virtually anyone can transform an idea into reality, the potential for progress multiplies. Yet those already adept at coding will see their capabilities grow exponentially, fuelling a surge in hyperinnovation. At the same time, these rapid advances are likely to heighten resistance from those wary of - or adversely affected by - swift technological change, further polarising cultural discourse. By acknowledging and respecting such counter-movements, we can pursue a future that balances human values with technological progress, ensuring that innovation serves us all.



Beyond Three Dimensions

Thomas S. Sibley

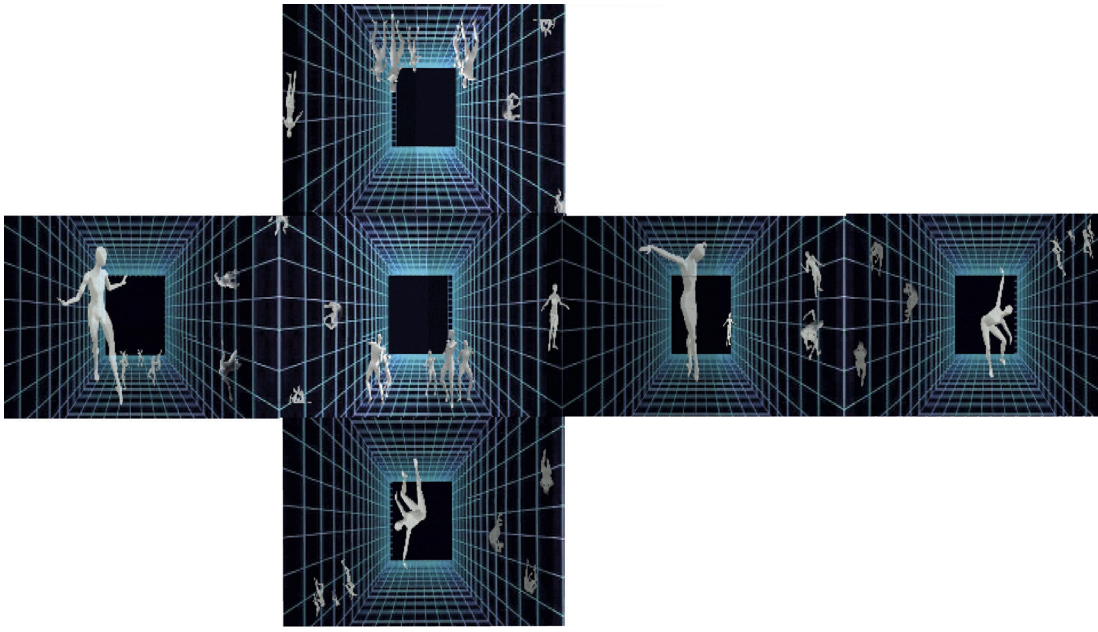
DOI: 10.4324/9781003726500-15

While studying holography and optics in college, I became fascinated with the idea of spatially projecting holograms and virtual reality, and have spent much of my life developing new technologies to solve these and other light puzzles. Through trial and equal degrees of success and error, I have developed several methods for projecting volume and binocular depth in extended reality (XR) environments.

Recently, I developed an auto-stereoscopic technology that uses aspects of eye/brain physiology and 4D geometry to project an immersive, hybrid reality. These simulations are visible to the naked eye without 3D glasses or headsets, bringing XR into the real world. Using a combination of AI, XR, auto stereoscopy, linear perspective and higher dimensional geometry, this immersive, hybrid reality is evolving into a new form of hyperspatial theatre known as the Tesseract Project.

Two Dimensions to Three

When predicting the future of 3D visualisation, it's useful to consider the history of linear perspective and its impact on art and science. In 1400 AD, the first mathematical approach to linear perspective was invented by Filippo Brunelleschi, 'using surveying techniques to plot the elevations, measurements "from base to base" and simple calculations based on triangulation'.[1]

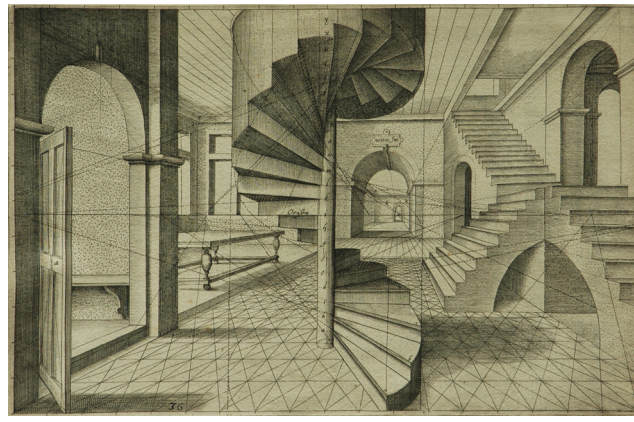


Tesseract Unfolded. Hyperspatial theatre immersion by T. Sibley and Cyberräuber. Lidar scans and 3D models of Badischen Staatsballet by Cyberräuber.

Artist Giotto soon followed with his geometric system of perspective and his paintings reflect a system that he described as 'those lines and planes situated above eye level should appear to decline downwards as they move away from the spectator; those below eye level should incline upwards; those to the left should incline inward to the right, those to the right should incline inwards to the left; there should be some sense of the horizontal division and the vertical division which mark the boundaries between the zones'.[2]

This new scientific approach to accurate perspective would have a major influence on art, science and culture for centuries to come. Drawing and painting with measured linear perspective accurately reproduced depth of field and geometric volume using relative scale, forced perspective, vanishing points and spatial coordinates of the line. The third dimension could now be rendered on a two-dimensional canvas as an accurate approximation of reality, becoming a virtual mirror wherein objects and scenery were reflected as projections of the mind.

The idea that humans could mimic and mirror the 'divine' by creating a parallel world became widely accepted as linear perspective became the key to projecting 3D space on a 2D surface. This signalled a quantum leap forward in the representation of three-dimensional space and would influence artistic and scientific thought for centuries to come. Humans could now represent precise space through dimensional scale which in turn redefined their position, space and place in the universe.



'Perspective, that is to say, the very renowned art of the eye point of view inside or through looking, being on a plain wall, on a painting, or on a canvas, in which there are some buildings…' A print from *Perspectives*, written by Vredeman de Vries and published in the Hague in 1604.

Three Dimensions to Four

Fast forward to 1880 and contemporary thought starts to puzzle over higher spatial dimensions as a new mathematical language evolves to describe them. Just as Renaissance thinkers jumped from the second to the third dimension, Victorians now began to ponder the possibility of a fourth dimension and higher.

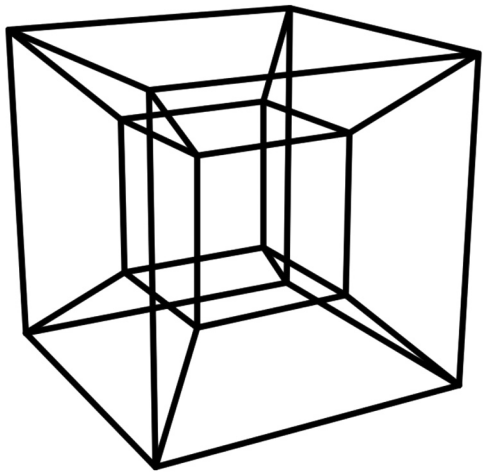
Three early pioneers to define this hyperspatial realm were mathematician Charles Howard Hinton, who coined the term 'tesseract', Anglican priest and theologian Edwin Abbott Abbott, author of *Flatland: A Romance of Many Dimensions*, and Alicia Boole Stott, a contemporary mathematician who took their ideas to the next level.[3]

Hinton published *The Fourth Dimension* in 1904 based on mathematical concepts he had been developing since 1880 that 'sought to establish an additional spatial dimension which he diagrammed as the "Tesseract"'.[4]

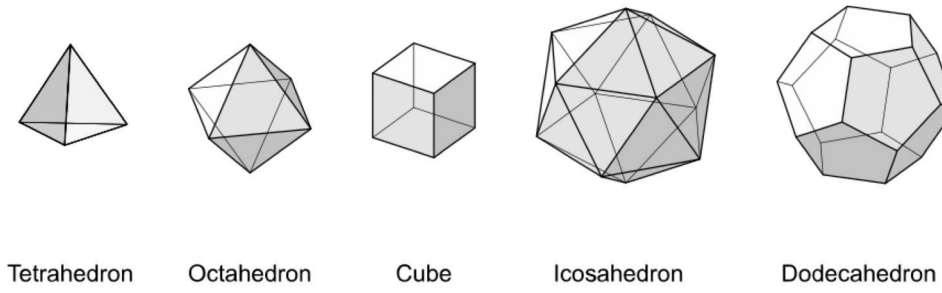
'In Greek, "τεσσάρα", meaning "four", transliterates to "tessara" and -act likely comes from "ακτίνες" meaning "rays"; so Hinton's use suggests the four rays from each vertex, exhibited in a tesseract and neatly encodes the idea of "four" into his four-dimensional polytope.' A polytope is a geometric object with flat faces or sides. However, in Latin, 'tessera' can also mean 'cube'.[5]

Hinton would move on to the next dimension at the early age of 54, but not before popularising the concept of higher dimensions and a means for visualising them to the general public. Interest in higher dimensions was high at the beginning of the 20th century, seeming to reinforce spiritualist concepts such as the spirit world and ghosts.

Hinton's publisher did not shy away from this supernatural association and titled Hinton's 1884 pamphlet, *Ghosts Explained*.[6] Hinton's collaborator, Alicia Boole Stott, would continue discovering higher dimensions, 'developing four-dimensional analogues of the cube, tetrahedron, octahedron and the dodecahedron, and built models of the four-dimensional equivalents of the Platonic solids'.[7] She coined the term polytope to 'describe a four-dimensional convex solid'.[8] It has come to mean any polyhedron or geometric solid with flat sides or faces.



Tesseract or hypercube.



Platonic solids provided as examples of regular polytopes.

Abbott's *Flatland* is a Victorian social satire set in a geometry-based class system. In *Flatland*, two-dimensional beings start to observe higher dimensional forms as 3D slices start falling through their 2D world.

By drawing a parallel between the spatial enlightenment of the second dimension as it experiences the third, Abbott suggested a means by which the fourth and higher dimensions might be visualised by inhabitants of the third. The fourth dimension as a spatial construct would later be replaced by Einstein in the popular and scientific lexicon as time.



Cover of *Flatland: A Romance of Many Dimensions* by Edward Abott Abbott.

Envisioning Four Dimensions

'The tesseract is a four-dimensional hypercube. It is the four-dimensional analogue of a cube, just as a cube is the three-dimensional analogue of a square.'[9]

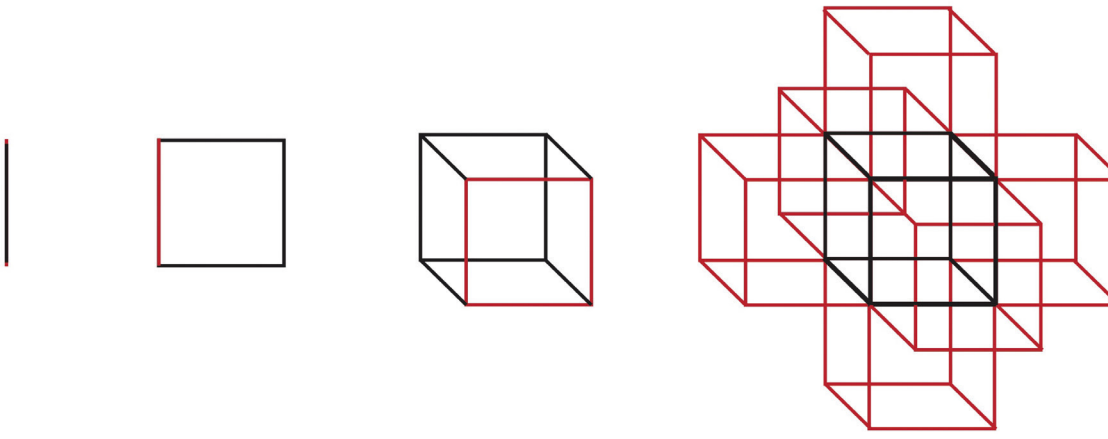
Another way to envision higher dimensions is to start with a point in space. The point exists in the zero dimension until it is connected with another point to become a line or the first dimension. When two equal horizontal lines are connected at the ends to two equal vertical lines, they become a square and the second dimension. Here they occupy two dimensions of space on the X and the Y axes (width and length). When two squares are connected by two additional squares on the Z axis (depth), they combine into the X, Y, and Z axes and become the third dimension.

Upward progression through dimensions can also be understood by seeing each dimension as being bound on its sides by the previous dimension. The line is bound on either side by points (1D). The shape is bound on its sides by lines (2D). The solid is bound on its sides by shapes (3D).

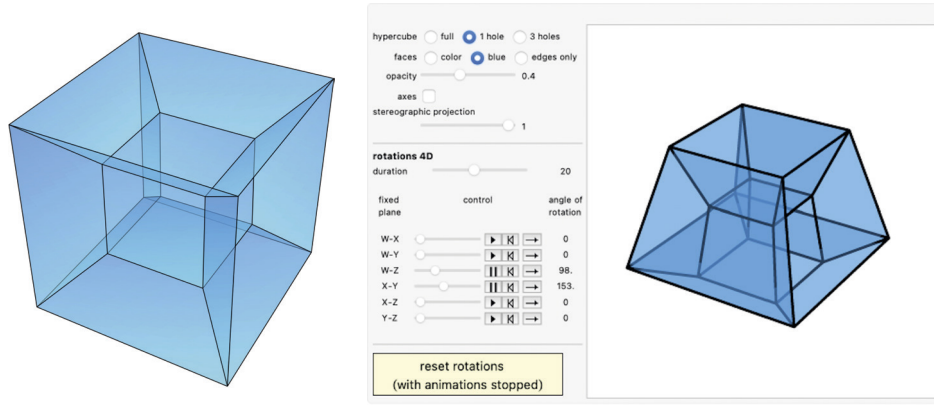
Taking the dimensional progression one step further from *Flatland* to 4D, the 4D tesseract, or hypercube as it is now known, is bound on all of its hypersurfaces by 3D cubes.

A hypercube can only be seen as an analogue or shadow of itself in the third dimension and is generally depicted in wireframe to permit multisided viewing. Computer model visualisations assist with conceptualising the higher dimensional structure of a hypercube in 4D by rotating and inverting the 4D tesseract upon itself in 3D. These animated visualisations permit the viewer to observe the tesseract from multiple angles and sides as it moves through time.

While the tesseract or hypercube forms the underlying architecture of the author's mixed reality system, autostereoscopy (binocular depth visible to the naked eye) is also necessary to produce a convincing hybrid world.



Upward progression of dimensions bound on all sides by the prior dimension.

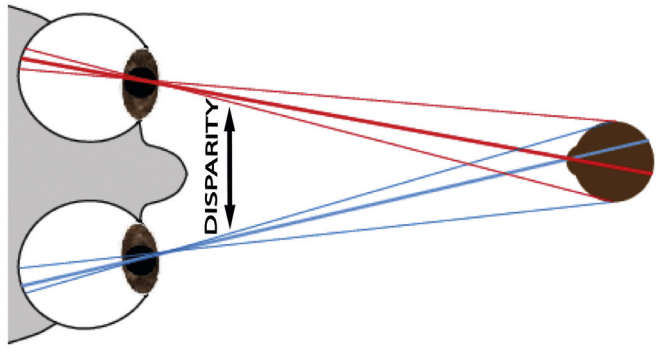


Tesseract or hypercube in two states.

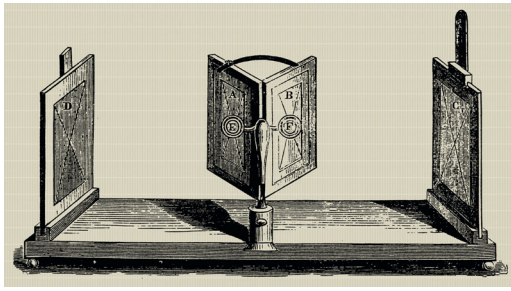
Stereoscopy

During the first half of the 19th century, underlying biological mechanisms for binocular and stereoscopic vision were becoming understood. Human vision is based on the eyes seeing the world at two slightly different angles. The brain converts disparities between the eyes into depth. The eyes combine two slightly different 2D views into a 3D object and/or environment in the brain.[10] This function of eye/brain physiology enables us to perceive, manipulate and navigate 3D objects and spatial coordinates.

The 19th century not only redefined dimensional horizons, but also expanded the applied sciences through advances in optics and the invention of photography. By exploiting the underlying stereoscopic mechanisms that convert 2D to 3D information in the brain, Victorian inventors discovered an optical means by which to simulate binocular vision. Charles Wheatstone invented the Stereoscope in 1838. The first description of his stereoscope was delivered to the Royal Society in his article, 'Contributions to the Physiology of Vision.—Part the First. On some remarkable, and hitherto unobserved, Phenomena of Binocular Vision'.[11]



Disparity in the binocular field.



Charles Wheatstone's mirror stereoscope.

Wheatstone's stereoscope used two mirrors (later replaced by binocular optics) positioned at a 45-degree angles to the viewer's eyes, reflecting images to each side that converged and conveyed both depth and volume. Hugely popular during the 19th century, stereoscopes permitted users to visit and experience remote people, places, and objects inaccessible due to the mobility constraints of the time.

One might argue that stereoscopes were the first version of virtual reality as they relied upon similar optics to modern headsets and enabled their audience to visit and experience remote locations and immersive realms of the imagination.

The Information Age

Today we exist in the ever-evolving Information Age where generative AI is approximating the technological singularity. AI has become an increasingly important parallel mind that assists humans with projects and problems beyond our cognitive scope. Computers with ever-increasing processing power enable humans to immerse themselves in supercharged virtual worlds, pushing the temporal and spatial envelope of reality, approximating and in some ways superseding the original.

However, it could be argued that the Information Age, with its virtual experiences, has come at a significant cost. Studies show that our psychological immersion in mobile devices and social networks has induced a host of anxiety disorders and other digital pathologies.

The internet has become an electronic mirror, reflecting both the beauty and ugliness, the best and worst aspects of human society. The technology that made the majority of human knowledge accessible instantly, anywhere, conjures equally powerful visions of wonder and woe.

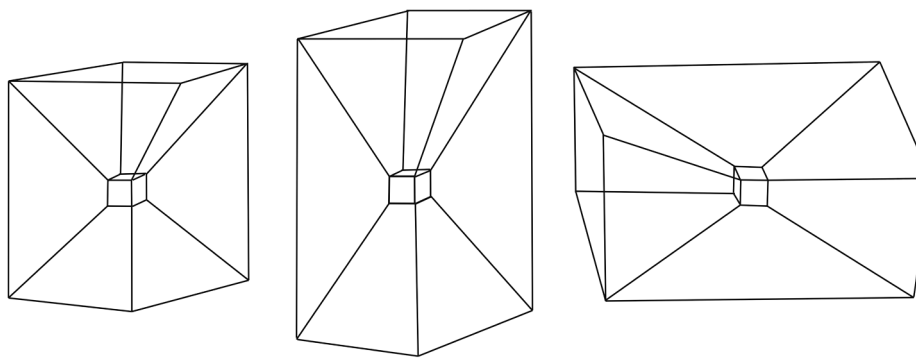
Virtual reality (VR), with its full immersion into other worlds, can be isolating and dehumanising. It brings the viewer into the virtual world through headsets as a solitary, confined observer and participant, separating people from reality. However, our evolving relationship with mixed reality is leading us towards hybrid models that blend the virtual and physical by immersing us in real-world simulations, increasing connectivity and the potential for collective experiences.

The Tesseract Project

Building upon my 3D imaging patent of 2023, The Tesseract Project introduces one such hybrid model in the form of a hyperspatial theater.[12] By wrapping an XR outer cube around a real space (Rn) inner cube or core and converting it to a hyperspatial environment, scale and other parameters become variable. Viewers within the inner cube are surrounded on all sides by an autostereoscopic immersion with simulated depth and geometric volume. The XR outer cube exists as a scalable extension of the Rn inner cube and an example of a projected relative or transcendent dimension.

The hyperspatial topography within this hybrid world is based on reverse engineering key concepts of Wheatstone's stereoscope into a spatial or autostereoscopic projector. This spatial projection technology enables the brain to register geometric volume and binocular depth of field, visible to the naked eye from multiple viewing zones. This and other emerging technologies, with AI at the forefront, have the potential to merge virtual and physical reality without the hitherto requisite VR headsets or 3D glasses. Rather than immersing the viewer inside the computer, they bring the computer simulation into the real world.

Projecting this autostereoscopic information in six directions creates a hyperspatial environment that encapsulates the viewer in a dynamic immersion. The tesseract, once a mental exercise for conceptualising the fourth dimension, has now become a means for projecting higher dimensions onto the real world.



Cross sections of XR architecture in three states: normal, vertical expansion and horizontal expansion.

Hyperdimensional Cities

In 2024, GPT-4 (a multimodal AI model approximating human reasoning) was asked if humans might one day be able to experience higher dimensions, and GPT-4's reply came in the form of a question.[13] It wanted to know if humans could augment and enhance their cognitive abilities. When answered in the affirmative, it offered several suggestions as to how humans might overcome their present cognitive limitations and comprehend hyperdimensional concepts or '3D conceptual shadows', such as the hypercube.[14]

While AI suggested several enhancements exceeding the scope of this discussion, the specific enhancement we're concerned with here is 'multi-layered reality perception'.[15] In the words of AI: 'in this multi-layered reality perception, each layer represents a different set of rules, properties, and interactions, some of which may be vastly different from your familiar physical laws. The layers are not separate but rather interconnected and interdependent, forming a complex web of causality and influence. Within this framework, conscious beings may have the ability to navigate and interact with these multiple layers of reality, leading to novel experiences and insights that challenge our traditional notions of existence, causality, and identity.'[16]

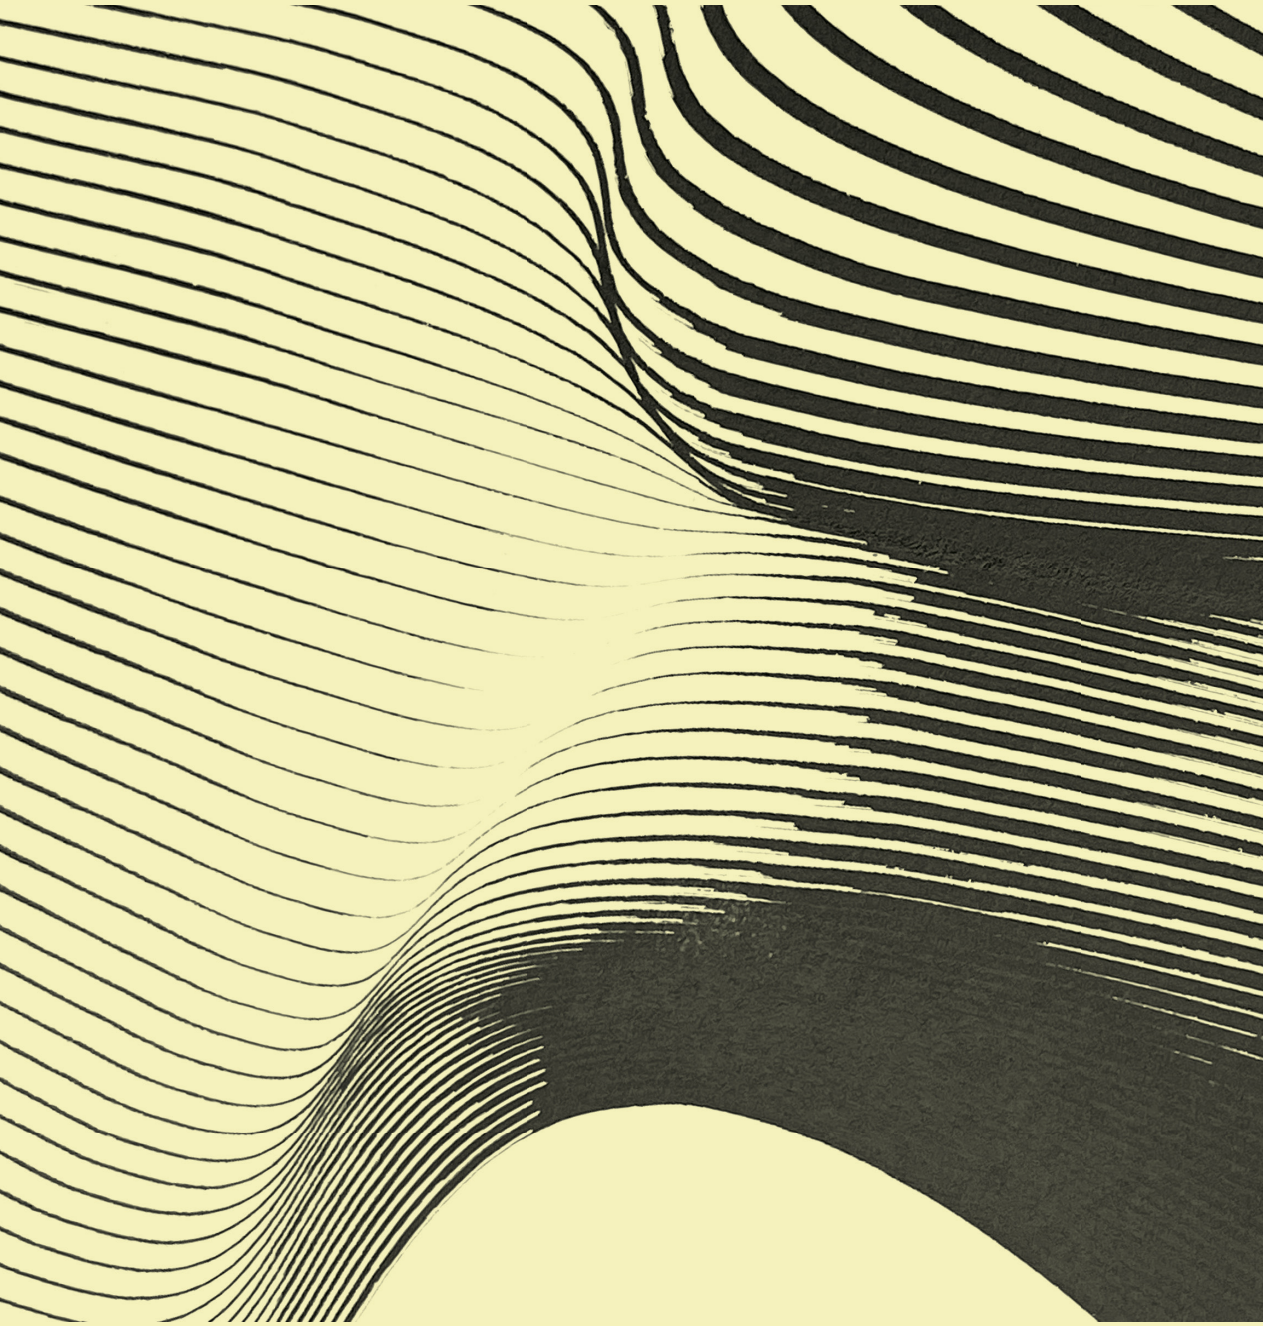
The future of hyperspatial architecture is based on this 'multi-layered reality' and the internal and external enhancement or augmentation of human perception.[17] This hybrid realm will evolve through the synthesis of physical and metaspatial dimensions. Existing between real and extended reality, variable in scale and style, these new worlds will exceed the structural limitations of reality.

Generative models that once existed as compressed data inside the computer will be mapped upon the human neural network and perceived reality as variable data points and assembled as higher dimensional structures and environments, dissolving and reassembling as new forms against a 360-degree event horizon.

This hyperspatial reality shall evolve as the inevitable intersection of mind and matter. Generated from the synthesis of human and artificial intelligence, hyperdimensional cities shall rise from the neural architectures of the mind.

Notes

1. M. Kemp, *The Science of Art: Optical Themes in Western Art from Brunelleschi to Seurat* (New Haven CN, 1990), pp. 9–13.
2. Ibid.
3. J. Crabb, 'Notes on the Fourth Dimension', *The Public Domain Review*, 2015, p.15.
4. Ibid.
5. Ibid.
6. M. Blacklock, 'The Four Dimensional Life of Mathematician Charles Howard Hinton', *BBC Science Focus*, (2020), https://www.sciencefocus.com/science/the-four-dimensional-life-of-mathematician-charles-howard-hinton, accessed 27 January 2024.
7. Ibid.
8. J. J. O'Connor & E. F. Robertson, 'Alicia Boole Stott', *MT Math Tutor*, (2014), https://mathshistory.st-andrews.ac.uk/Biographies/Stott/, accessed 3 March 2025.
9. A definition of Tesseract can be found on Wikipedia, https://en.wikipedia.org/w/index.php?title=Tesseract&oldid=1199603873, accessed 27 January 2024.
10. F. Gonzalez & R. Perez, 'Neural mechanisms underlying stereoscopic vision', *Progress in Neurobiology*, 55(3) (1998), [doi: 1016/s0301-0082(98)00012-4], pp. 191–224.
11. C. Wheatstone, 'Contributions to the Physiology of Vision.—Part the First. On some remarkable, and hitherto unobserved, Phenomena of Binocular Vision', *Philosophical Transactions of the Royal Society of London*, vol 128 (1838), pp. 371–94.
12. T. Sibley, Projection System and Method for Three-Dimensional Images, U.S. Patent and Trademark Office, U.S. Patent № 11,595,628 (2023).
13. P. Van Treuren, 'Artificial Intelligence and "4D" Philosophy', *Stoic Simple*, https://www.stoicsimple.com/artificial-intelligence-and-4d-philosophy/, accessed 4 March 2025.
14. Ibid.
15. Ibid.
16. Ibid.
17. Ibid.



Robotic Mark-Making

ChatGPT-4 + Python + Grasshopper + KUKA|prc

Sara Codarin

DOI: 10.4324/9781003726500-16

JOURNEY

Digital fabrication occurs through the process of translation, where digital models are converted into data streams that machines can interpret and execute. Despite the automated nature of this process, there are opportunities where robotic dynamics can intersect with human creativity and agency. This workflow showcases how the analogue process of painting is reconsidered through the lens of human-machine interoperability, using AI as an accelerator of results.

An industrial robotic arm is equipped with a 3D-printed actuator. This device holds a soft-brush pen programmed to draw ink lines on a target area within the bounds of a paper sheet. Instead of working on a flat bidimensional surface, the soft brush navigates a 2.5D space, transforming precise digital inputs into one-off ink marks. Dense linework can blend into ink patches, while interweaving patterns alter the legibility of each layering sequence.

The creation of the drawings is informed by mathematical equations sourced from interrogating large language model platforms (LLMs), including ChatGPT. This approach facilitates the generation of Python scripts, subsequently integrated into Grasshopper. This transition establishes a shift from purely geometric to mathematically driven instructions, incorporating control parameters directly within the scripts' functions.

Finally, a custom Python script component in Grasshopper is defined to fine-tune the digital-to-physical conversion. Choosing to work in the 2.5D space brings forth a deliberate embrace of anomalies and uncertainty in the outcome of the designed paths. This approach involves continuous interaction between the digital toolpath and the physical act of mark-making by the robot, which highlights the divergence between intention and unpredictability in the creative process.

Each outcome serves as a testament to the craft and dynamic interplay between digital instructions and physical execution, where two-dimensional marks carry the depth of three-dimensional implications.

WORKFLOW

STEP 1

Start by engaging ChatGPT to explore the form-finding of linework, focusing on identifying a mathematical function capable of producing Python-driven, 3D digital lines. Inspiration can be drawn from the physical description of natural phenomena or spatial representations of chaotic systems like the Strange Attractors. The prompt craft is minimal: 'generate a Gh Python script that creates 3d points (x,y,z) based on the Aizawa Attractor equation'.

```
import Rhino.Geometry as rg
import random

# Define the Lorenz attractor equation
def lorenz(x, y, z):

    sigma = 10
    rho = 28
    beta = 8 / 3

    dx = sigma * (y - x)
    dy = x * (rho - z) - y
    dz = x * y - beta * z

    return dx, dy, dz

# Set the number of points to generate
num_points = 10000

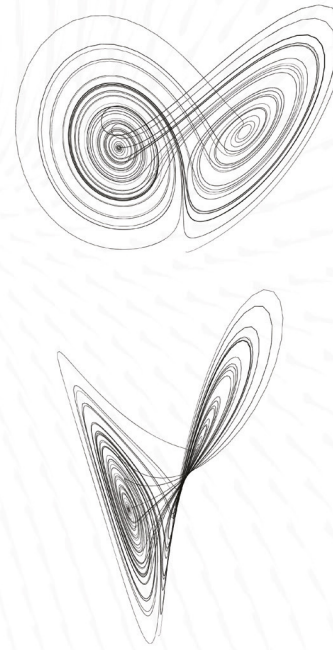
# Set the starting point
x, y, z = 0.1, 0, 0

# Generate the points
points = []

for i in range(num_points):

    dx, dy, dz = lorenz(x, y, z)
    x += dx * 0.01
    y += dy * 0.01
    z += dz * 0.01
    points.append(rg.Point3d(x, y, z))

# Export the points to Grasshopper
a = points
```



```
import Rhino.Geometry as rg

# Define the Aizawa attractor equation
def aizawa(x, y, z):
    a = 0.95
    b = 0.7
    c = 0.6
    d = 3.5
    e = 0.25
    f = 0.1

    dx = (z - b) * x - d * y
    dy = d * x + (z - b) * y
    dz = c + a * z - (z ** 3) / 3 - (x ** 2 + y
** 2) * (1 + e * z) + f * z * x ** 3

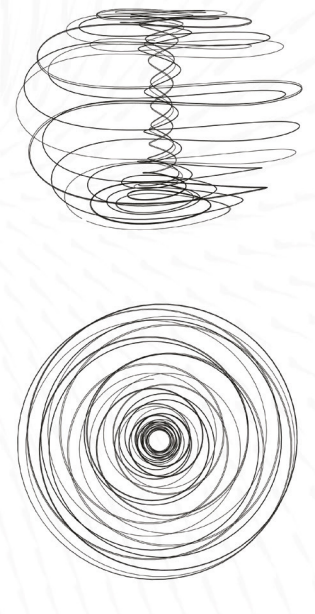
    return dx, dy, dz

# Set the number of points to generate
num_points = 10000

# Set the starting point
x, y, z = 0.1, 0, 0

# Generate the points
points = []
for i in range(num_points):
    dx, dy, dz = aizawa(x, y, z)
    x += dx * 0.01
    y += dy * 0.01
    z += dz * 0.01
    points.append(rg.Point3d(x, y, z))

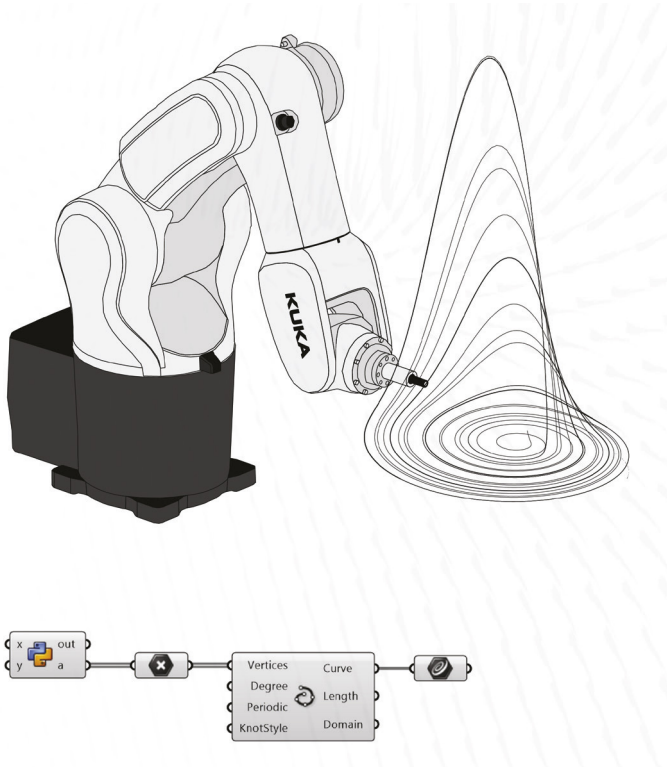
# Export the points to Grasshopper
a = points
```



(Top) Prompt-driven 3D curves describing the Lorenz Attractor mathematical equation.
(Bottom) Prompt-driven 3D curves describing the Aizawa Attractor mathematical equation.

STEP 2

Import the generated Python script into Grasshopper and use the component ‘Interpolate’ to visualise the 3D curves.



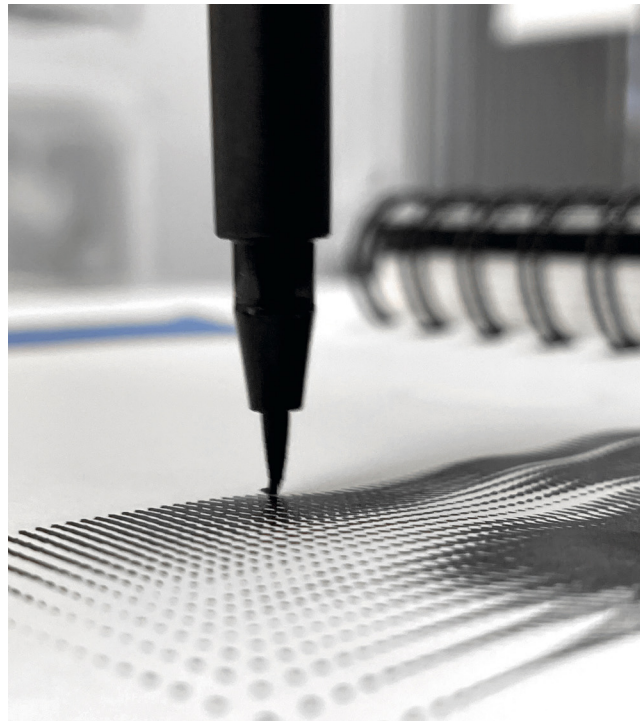
Manipulation of the AI-generated script in Grasshopper.

STEP 3

Use Grasshopper to frame the curves into a bounding box defined by the dimensions (x, y, 0.25mm). Here, x and y correspond to the dimensions of the intended drawing area, and the z value signifies the maximum depth to which the physical brush is allowed to press against the paper. Consequently, the soft brush moves through the 2.5D space, transforming the precise digital inputs into unique ink marks. These marks reflect the varied iterations, adjustments and external factors that shape the final outcomes, including ink saturation, paper texture, atmospheric humidity, the robot’s execution speed and the creative whims of the designers.



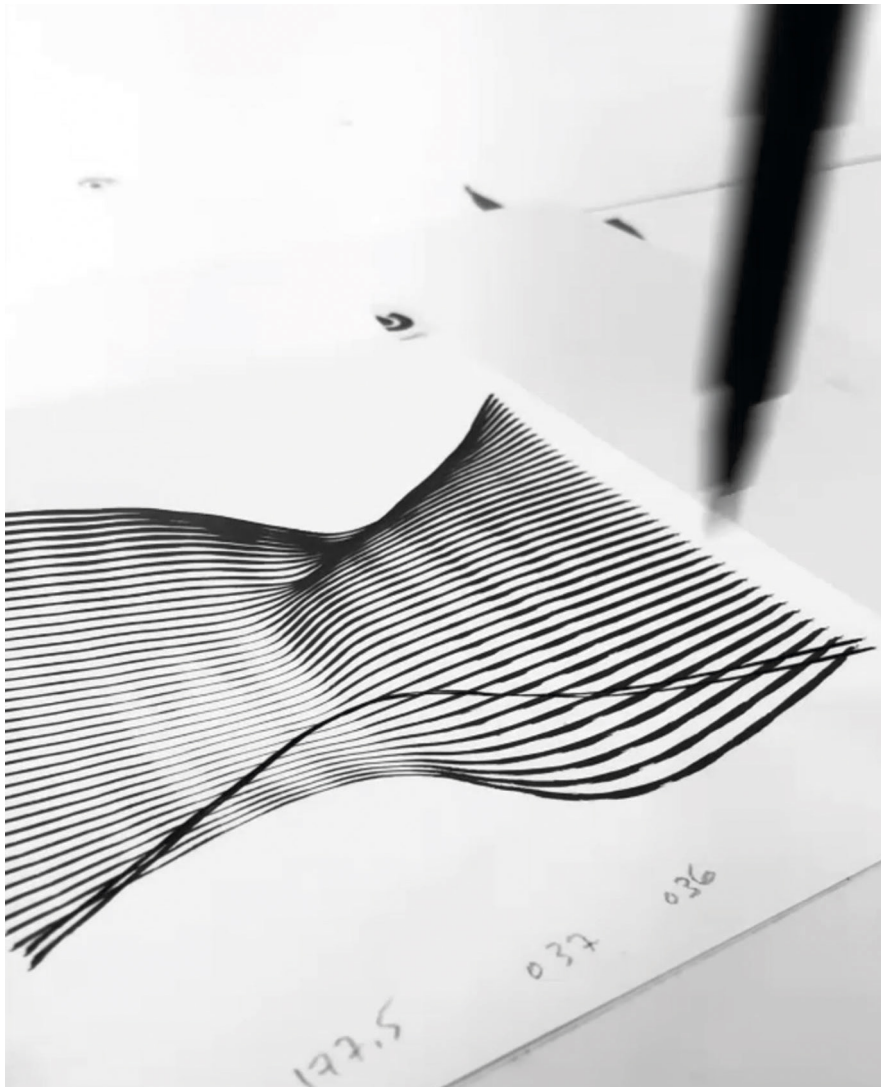
Transformation from 3D to 2.5D geometry.



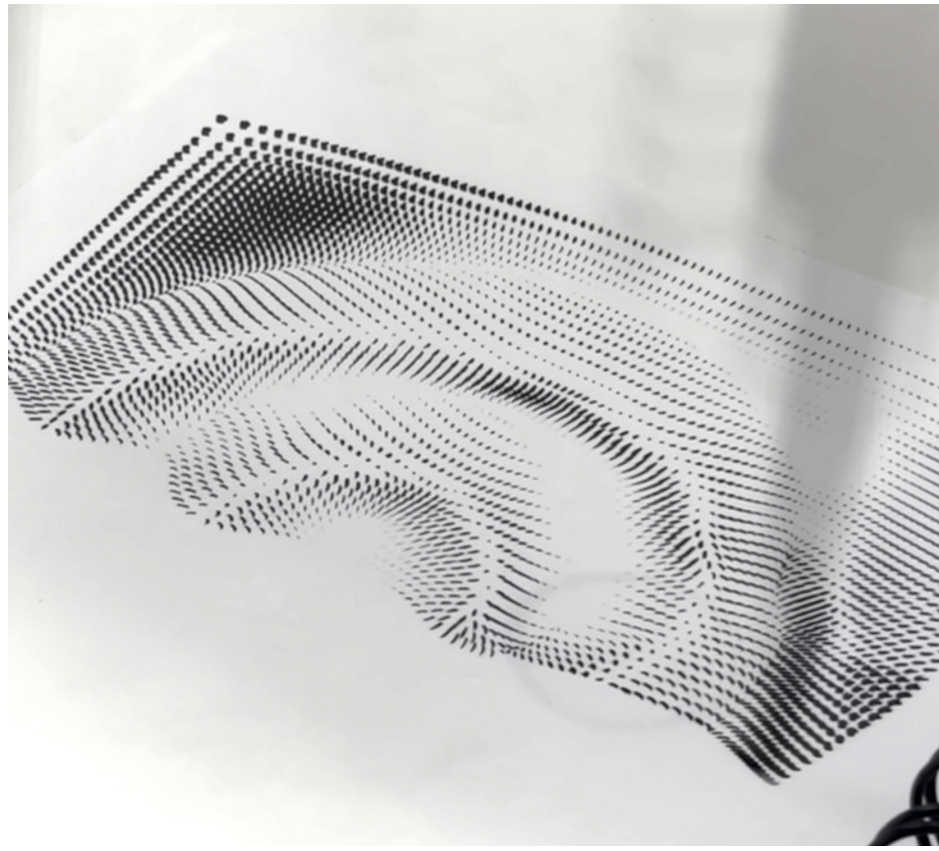
Robotically controlled soft-brush pen navigating the 2.5D space.

STEP 4

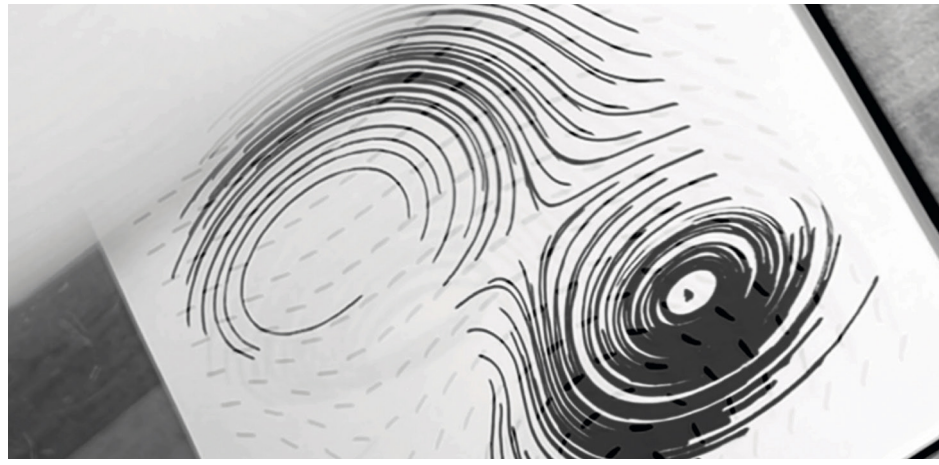
Calibrate the robot's tool centre point to align with the line nearest to the upper boundary of the drawing's bounding box. This calibration allows for achieving the desired thickness of the curves. Start by setting the initial calibration depth to match the length of the soft brush, with subsequent fine-tuning possible down to 1/10 of a millimetre. Once the initial curve is tested for accuracy, the robot is ready to proceed with executing the full drawing.



Definition of calibration lines within the AI-to-geometry-to-fabrication workflow.



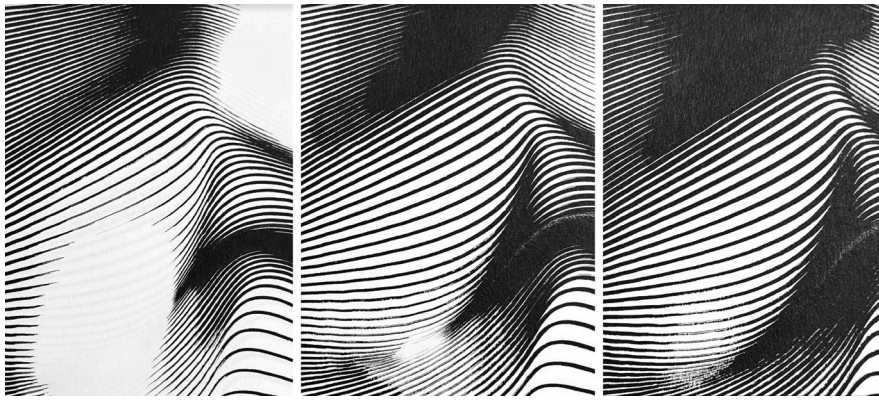
Robotic implementation of AI-generated 2.5D drawings.



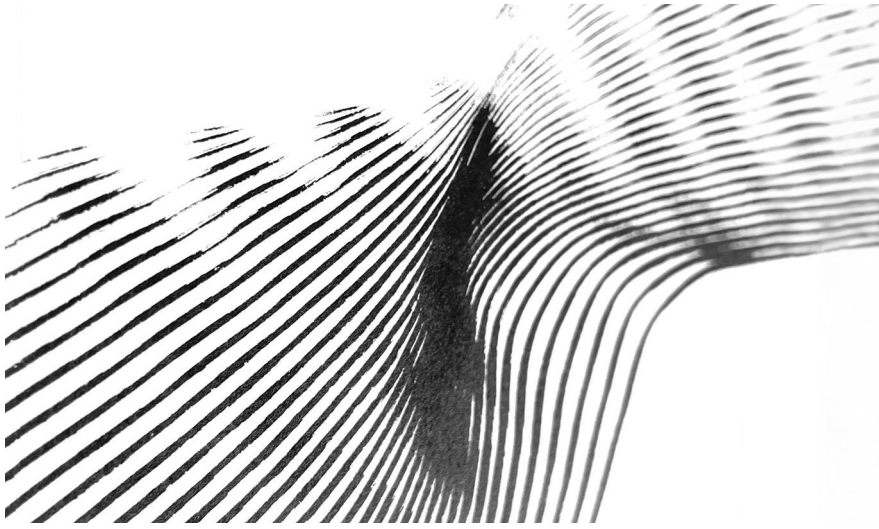
Outcome.

STEP 5

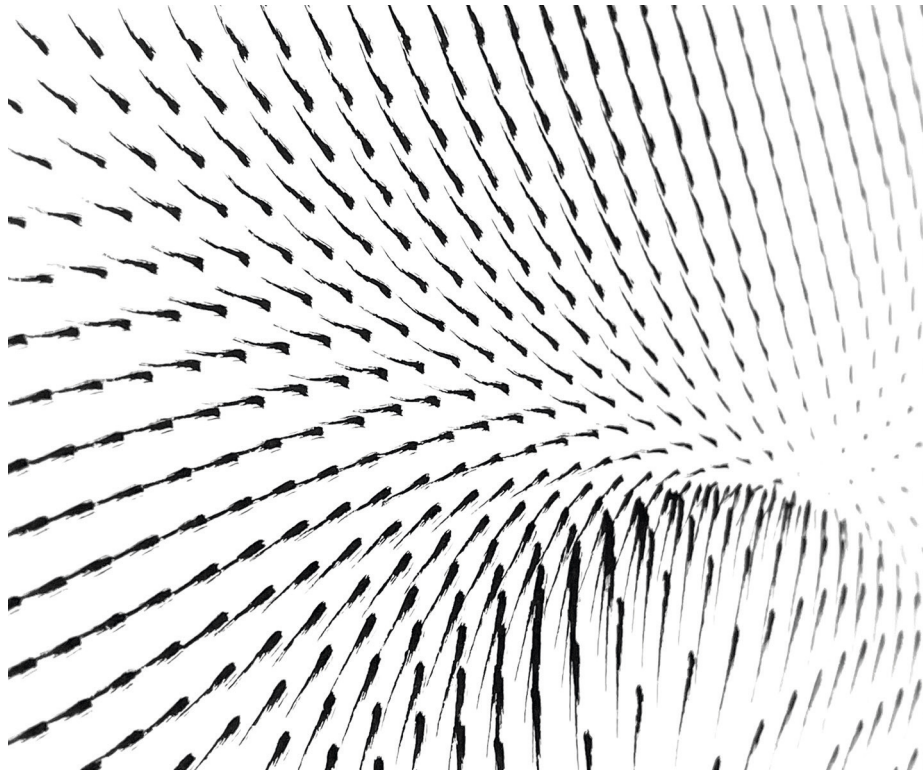
Employ varied tool configurations for multiple iterations, documenting the subtle distinctions across successive executions of the same drawing script. Examining the drawings more closely highlights the individual ink traces, drawing focus on the directionality of the toolpath programmed for their creation. This approach to design strategy considers AI-generated data, time, motion, actuators and ink on an equal footing as materials in the creative process.



Iterations obtained by adjusting hardware and software settings.



Detail of the 2.5D robotically implemented drawing.



Detail of the 2.5D robotically implemented drawing.



Detail of the 2.5D robotically implemented drawing.

TIPS

Embrace indeterminacy	View AI as a continuously evolving open process with endless possibilities. Observe glitches and transitions, exploring the spaces between dimensions. View outcomes as blends of successive visions rather than finalised products, leveraging AI's variability to enrich your creative workflows.
Bring humanistic context	Infuse your vision and cultural background into AI-generated outcomes, as AI lacks inherent cultural context or situated embodiment. Adopt a humanistic perspective to avoid a technocentric approach in using technology, drawing upon AI's data malleability to substantiate the meaning of creations derived from its contributions.
Establish your authorship with autonomous technology	Collaboratively use AI tools to engage collective intelligence built from datasets while taking responsibility for editing and curating the generated content. This mindset upholds your personal authorship and maintains the quality and integrity of your design work.

What does the future role of the architect or designer look like?

Architects and designers have long integrated advanced technologies to expand creative boundaries. Today, AI adoption marks a global shift, transforming design processes. The outdated human-machine divide is unproductive – designers now collaborate with AI, accelerating results through iteration and optimisation. This redefines their role, focusing on data-driven workflows and systems to tackle complex challenges. Italian architect and designer Aldo Rossi noted that architectural ideas hold value beyond physical form; AI amplifies this, emphasising processual thinking. Workflows become integral to a project's narrative, blending creativity, logic and beauty. The future of design lies in this fusion of human vision, technology and process as a creative force.



Data is the Answer, but What is the Question?

Hamza Shaikh

DOI: 10.4324/9781003726500-17

Many of us in the architecture, engineering and construction fields were left wondering why, more than three years after the generative AI boom of 2022, leading 3D-modelling platforms were yet to deliver significant AI-driven capabilities for spatial, programmatic and technical design. Despite the rapid advancements in generative technologies offered by large language models (LLMs) and AI training breakthroughs, years later, the latter stage technical workflows for building design remained relatively underdeveloped. Programs to generate entire detailed and custom floorplans, 3D-detailed façades and resolved structure solutions have not kept pace with the strides made in image generation innovation for concept level ideation. Studies such as a 2024 survey by the MIT Sloan School of Management, highlighting that only 4% of firms in the construction industry use AI compared to 12% in sectors like manufacturing, underscore this lag in adapting AI-driven innovations to technical workflows in building design. The question persists: Why is this so?

The answer lies in one critical, complex challenge – data.

While huge amounts can be done with image generation technology that relies on vast datasets of tagged, often publicly available images, the demands of data-rich 3D models are not as easy to meet. This is because 3D data typically requires intricate metadata to define spatial and material properties, making it far more complex than flat 2D images. Moreover, 3D models are often proprietary, stored within specialised software ecosystems, or created for specific projects, making them harder to access or standardise for large-scale generative training purposes. This puts architecture and engineering firms in a strong position to leverage all their design data for proprietary solutions, but doing so requires significant investment and risk.

For any generative tool to move beyond raster-based image outputs and produce rapid, data-embedded vectorised assets or 3D configurations suitable for building information modelling (BIM)-level design – whether it's generating custom floor layouts and furniture for a residential tower, a rationalised parametric 3D model of a completely new bridge design or custom interior material IDs and finishes – it requires training on extensive and highly specialised datasets. Fortunately for some organisations who have spent decades organising and engineering their project data meticulously, shifting into this next era of data-driven generative design will be more seamless.

For most design firms, past 3D models and project data are usually stored primarily for legal protection and record-keeping for accreditation purposes, however, firms now may find themselves considering how to leverage their past data for AI training capabilities. Having large amounts of data is not the answer if it's a complete mess, but having 'good' data is critical. To improve data quality, firms can start by implementing robust data management systems, standardising file formats and ensuring consistent metadata tagging across projects. Additionally, hiring data strategists and engineers, investing in training programs for staff to maintain and curate datasets effectively, as well as periodically auditing and cleaning historical data, will help kickstart the process. Nonetheless, those who prepared for the data-driven future years ago will bear the fruits going forward.

Transitioning to Generative 3D

Of course, saying data is the answer is somewhat a given – data is an incredibly broad term. To understand the intentionality needed within the architecture and construction



Depiction of raster-based mesh geometry.

world with regards to the future of data-driven technology, we first need to understand the importance of the type of data most critical to the industry – 3D data.

In an era dominated by generative 2D advancements, the emergence of generative 3D technologies is set to reshape industries ranging from film, gaming, design and construction. The design and construction sectors however, with their more complex demands, require far more data-rich foundations and arguably a different approach completely. In gaming, for example, generative 3D has already transformed workflows, enabling the creation of expansive world-building assets with unprecedented speed and efficiency. Titles like *Cyberpunk 2077* and platforms such as Unreal Engine 5 illustrate how tools like procedural asset generation streamline creative processes. However, these advancements largely rely on raster-based methods, including neural radiance fields (NeRFs) and point cloud meshes, which lack the data richness and precision required for applications like BIM. NeRFs and point cloud meshes are primarily focused on visual representation and spatial approximation rather than delivering the granular, multidimensional data needed for BIM. They often lack integration with real-world metrics and fail to incorporate essential layers of information, such as structural integrity, regulatory compliance and material specifications, which are critical for construction and architectural applications.

Rasterised 3D assets suffice for screen-based outputs in gaming and film, but they fall short in construction, where physicality, precision and data interoperability are paramount. Unlike gaming assets, BIM-level data must incorporate layers of information – structural, regulatory, material and environmental – aligned to real-world metrics. For instance, imagine two scenarios – one generating a virtual urban environment for cinematic storytelling and the other for urban planning and construction. In cinema, the focus is on sensory immersion; in urban planning, the environment must incorporate real-world physics and therefore viable structures, layouts that conform to zoning requirements, building regulations, community

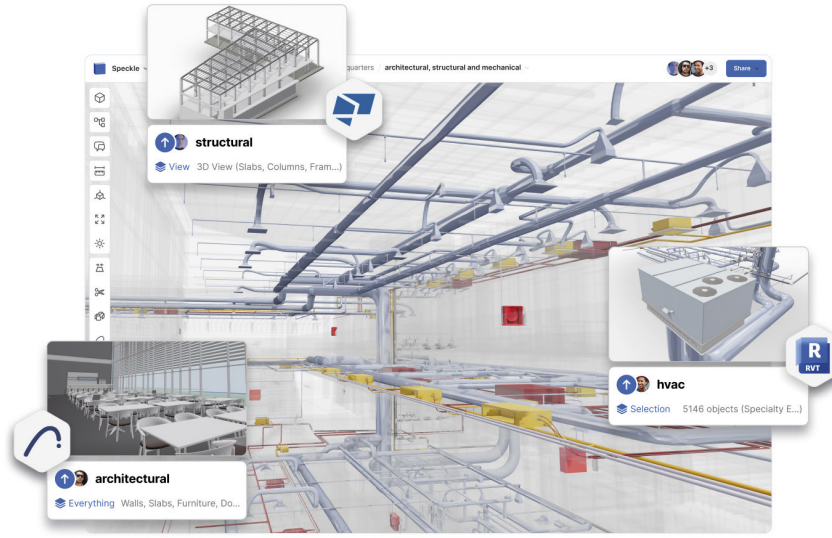
guidance and sustainability standards. If architects want to continue to realise their visions in the real world with the assistance of generative AI, then the future technologies require vast vectorised, multidimensional and dynamic datasets.

To further understand this need, let us consider the demands in a future scenario for a well-known 3D BIM software, Revit. Imagine being able to ask Revit to generate an entirely detailed and complete building for a five-storey mixed-use office building tailored to your firm's unique architectural style as well as extensive custom design specifications, situated on a specific plot in London. For such a request to be realised, the generative functions would require access to multiple rich and vast datasets from various sources: the firm's unique design data repository, London's planning and building regulations, supplementary design guidance documents, site-specific environmental data, key precedents, contextual vernacular design data and perhaps even a contextual qualitative dataset capturing the community's preferences. Beyond these, it would also need to incorporate structural engineering requirements and do this while driving all decision-making with sustainability targets. These are not abstract requirements; they represent the reality of the built environment – a complex interplay of tangible and intangible factors that influence every design decision. Achieving a scenario where such functions and data can be called upon and inferred within a design process presents significant challenges, yet it appears to be the inevitable end goal given the rapid and exponential advancements in generative AI development we are witnessing today.

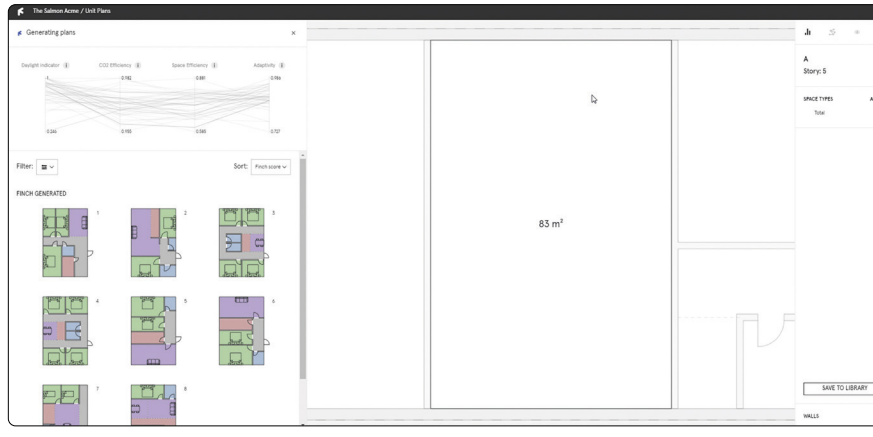
To keep up, organisations must first carefully define, or perhaps redefine, their data strategies and overall mission before meaningful work can be done with generative AI integration. Is the mission to provide real-time data insights from vast resources at every level of the design process to make the best-informed design decisions? Or is it to look inward at efficiency and speed by creating comprehensive data frameworks and protocols based on past work for future design decisions? In any case, having a refined overarching data mission allows organisations to begin creating effective pipelines, which is crucial to ensuring AI does not dictate the trajectory. This is because AI-enablement should be a facilitator of renewed goals as opposed to the goal itself. Furthermore, training AI models and developing application solutions that fit into this wider data strategy requires substantial investment and cross-industry collaboration.

Due to the uncertainty and rapid rate of development of AI technologies, emerging startups are making bold steps into this space. Companies like Speckle, Differential and Finch3D are pioneering generative design platforms albeit in very different ways. While currently still in their development phases, they demonstrate the immense potential for tailored, data-rich 3D design tools. For instance, Differential's Rhino M8 - Multi AI Aggregator plugin allows users to transform natural language prompts into 3D models, showcasing its ability to generate entire building masses with precise parameters, such as 'create a 40x40 tower, 30 levels, 4m level height, with 2-degree rotation'. Finch3D leverages graph-based and generative technologies to automate floor planning, adapt room layouts responsively, and perform live spatial analyses. Speckle on the other hand, creates a unified 3D model interface for seamless and customisable multiplatform integration. As an open-source platform, it facilitates data interoperability and collaboration across disciplines, allowing architects,

engineers and contractors to integrate data from different sources into a single customisable platform. By doing this, it enables real-time updates and feedback. In an AI-powered future, platforms like Speckle could serve as the backbone for integrating generative AI tools into architectural workflows, providing firms with the flexibility to plug in their own data libraries and tailor the program to their specific needs. These examples underscore how startups are driving innovation, challenging established software providers to adapt and placing the architectural, engineering and construction (AEC) industry at the forefront of a technological revolution. This period of shifting sands is redefining how organisations approach design, construction and the ownership of design and documentation processes.



Speckle platform UI showing collaborative multiplatform unified interface.



Finch3D platform UI showing real-time generative floor plans.

The Challenges and Opportunities of Being Data-driven for AI

The crux of the challenge is not just the volume of data required, but the need to incorporate the immensely wide-ranging datasets and format types from consultants and stakeholders. For example, datasets may include BIM models from architects, geographic information system (GIS) data from urban planners, material specifications from manufacturers and compliance documents from regulatory bodies. These datasets often come in disparate formats such as .ifc, .dwg, .csv or proprietary file types, adding significant complexity to the integration process. Regulatory codes, for example, differ not only between countries but also between municipalities. Material properties, construction methods and environmental conditions vary widely depending on geographic and economic contexts. All of this makes the creation of interoperable generative architectural datasets a monumental task. Each of these data sources within a pipeline will have their own file formats, adding further layers of complexity to an already formidable challenge.

This underscores the renewed importance of universal file formats and environments, which, while not a new concept – Industry Foundation Classes (IFC), for instance, have been around for decades – are gaining fresh relevance due to the expansive data landscape introduced by AI. For instance, IFC formats enable seamless data exchange across different software platforms, reducing errors during collaboration and ensuring consistency in complex workflows. This interoperability is particularly beneficial in AI-driven projects where multiple stakeholders rely on accurate and comprehensive datasets to optimise decision-making. The ability to integrate diverse datasets seamlessly will not only enhance collaboration but also address the growing complexity of architectural and construction processes. Moreover, this push towards interoperability is likely to be embraced by planning authorities and legislative bodies, as it brings increased clarity and traceability to decision-making. Initiatives like the UK's 'Golden Thread' policy highlight the demand for transparent, data-driven workflows that ensure safety, compliance and accountability throughout the building life cycle.

In this new landscape, established software providers must adapt their roles in a generative, data-driven future and have already begun taking steps. Autodesk, for instance, has introduced AI-powered tools such as Construction IQ, which analyses construction data to identify risks and improve project outcomes. Automated symbol detection and photo autotags in Autodesk Build reduce manual input and Forma produces AI-based embodied carbon analysis for early-stage design. Similarly, Trimble has embraced AI innovations like AutoShape for rapid geometry creation, AI-enhanced search features in 3D Warehouse, and an array of diffusion-powered features within their SketchUp app. Established providers are slowly but surely adapting their platforms to integrate advanced AI capabilities, meeting the evolving demands of efficiency and sustainability in the design and construction industries. However, given their well-entrenched positions and the industry's current reliance on their core functions, these companies are unlikely to pivot drastically from their fundamental offerings.

For all software providers, there is a delicate balance in listening to hot industry trends and short-term expectations from online 'noise' while staying true to foundational capabilities. Rhino, for instance, already excels in its plugin and

interoperability features, allowing other developers to integrate tools directly through systems like Grasshopper. This modularity has made Rhino a popular choice among computational designers for decades, enabling customised workflows that adapt to specific project needs. Given the evolution of emerging players in the generative 3D space, established companies may even strategically acquire startups specialising in generative AI and 3D data-driven innovations. This approach would allow firms like Autodesk and McNeel to incorporate cutting-edge capabilities while preserving the trusted reliability that their users rely on.

Determining a renewed data strategy, the age of AI is a complex but vital challenge. The future depends on how industries collaboratively adapt together, but is driven by the bold steps taken by early movers. Companies like Adobe and NVIDIA demonstrate how AI can transform value propositions and set precedent for others to follow. It also highlights the importance of adaptable business models in aligning with AI integration. AEC organisations are likely to diversify their approaches in response to AI integration in different ways. Larger firms may focus more on proprietary, data-driven AI systems to efficiently manage large-scale construction environments, leveraging bespoke data insights and more automated workflows. Meanwhile, medium and smaller firms could embrace more commercial AI tools and generative 2D creative concept design, expanding their offerings to include immersive and interactive virtual environments. This evolution could blur the boundaries across industries like design, film and gaming, creating new opportunities for collaboration and innovation. The future of design lies with those who combine generative creativity and data-driven 3D technologies, adapting to the rapid advancements in construction and manufacturing with precision and purpose.



Depiction of future robotic construction technologies assisting building material manufacturing and assembly.



AI and the Climate Challenge: The Data Centre Paradox

Aleksandar Bajic & Hamza Shaikh

DOI: 10.4324/9781003726500-18

In our discussions with experts in environmental, social and governance (ESG) frameworks and data centre facility management and design, the conversation often circles back to one pressing question: How can AI with its immense energy demands help to improve our environmental future? It's a query that must resist simple answers and 'greenwashed' thinking, but through these exchanges, several insights have emerged as potential paths towards a positive future for our planet.

Firstly, it is worth accepting that AI is neither inherently good nor bad, but rather its impact lies in how we wield and integrate it into existing systems. Its potential for transformative innovation is undeniable, yet so too are its drawbacks. The scale of AI's resource consumption forces us to confront uncomfortable truths about energy production and environmental impact. However, this very pressure should catalyse more urgent and innovative solutions to climate challenges, accelerating the transition to cleaner energy systems and more sustainable practices.

Life cycle assessments (LCAs) are a cornerstone in understanding the environmental footprint of any system, and AI is no exception. The stark reality is that the resource demands of AI – particularly finite minerals and rare earths – exacerbate an already precarious situation of global resource depletion. This is further compounded by the insatiable energy and water requirements of data centres, which are proliferating to meet AI's growing appetite. These facilities – essential to AI's operation – have become flashpoints in a larger debate about energy production, sustainability and environmental ethics to reach global net-zero goals and potentially net-positive systems.

However, scepticism around the term net-positive is understandable – it's a lofty ideal often dismissed as impractical. Yet, net-positive doesn't have to mean zero resource consumption. Instead, it could signify a system where resources are used intelligently, waste is minimised to near-zero levels, and subsequently, outputs of resource usage positively contribute to regenerative cycles. AI, with its capacity to optimise processes and identify efficiencies, could be instrumental in realising this vision – a solution to its own problems.

Imagine a system where energy demands are offset by AI-driven efficiencies in production, waste management and energy recovery, creating a feedback loop that approximates a net-positive outcome. Or where mineral mining and extraction is mitigated drastically due to the re-use of modular and interchangeable hardware. NVIDIA's Blackwell GPUs, which offer significant power savings compared to traditional CPUs, promote end-to-end sustainable computing and recycling programs for its products.[1] The broader semiconductor industry is advancing circular manufacturing processes, aiming to reuse high-value components like CPUs and memory modules, and is supported by policies like the United States' CHIPS Act, which funds clean energy technologies for chip fabrication.[2] Startups like Recycleye are also leveraging AI-driven systems to improve material recovery and recycling efficiency.[3] These efforts, combined with advancements in sustainable manufacturing and recycling technologies, reflect a shift toward a circular economy in AI hardware while addressing the urgent need to curb mineral depletion and reduce e-waste. Yet, growing urgency and development are needed as demands continue to grow.

This begs the big question: Can AI shift from being a net-consumer to a net-contributor? While the outright cessation of energy demands may be unrealistic, the concept of cyclical energy production emerges as a real possibility and arguably

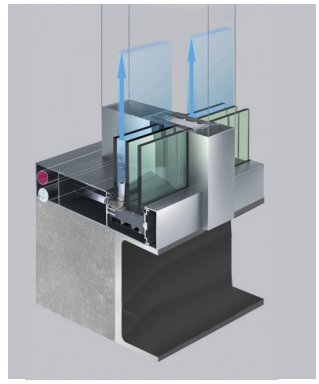
a non-negotiable necessity. In this context, a key example of a system in urgent need of rethinking in the design world are data centres, not just due to the operational carbon impact these buildings have, but the direct tie they have with the increasing demands of AI technology as organisations begin to require more server space. The immense heat generated by AI-driven data centres is often emitted as waste but shows huge potential to be harnessed and repurposed as a viable energy source. Novel methods are emerging that show the potential for capturing waste heat so that it can be converted into alternative forms of energy to power homes and businesses – one of these methods involving innovative architectural design is 'water-filled glass'.

Water-Filled Glass

Data centres are among the most energy-intensive buildings on the planet, with demand surging due to the growth of cloud computing and AI. Beyond their massive physical footprints, they pose a critical energy challenge, straining local power grids and consuming resources at unsustainable rates.

Efforts to address these challenges focus on three key areas: improving server technology, advancing cooling systems and recovering waste heat. While server upgrades yield incremental gains, cooling alone accounts for as much as 40% of existing data centre energy use, as they produce significant heat, much of which escapes into the atmosphere.[4] Capturing and reusing this heat presents a transformative opportunity to reduce energy consumption and offset resource use.

One promising approach could be to reimagine data centre façades with water-filled glass (WFG) technology. These panels have the capacity to harness the steady stream of low-grade heat generated by servers to produce thermal energy. This innovative system has significant potential to power nearby buildings and even entire communities. An adaptation of this product could even cultivate algae within the active water layer which presents a novel opportunity to synchronise warm water capture with biomass production, enhancing the energy reclamation of this architectural device. As detailed in the following passages, this concept forms the foundation of a proposal we could call urban low-energy networks (ULEN), which envisions data centres as sustainable, energy-positive anchors within urban ecosystems.

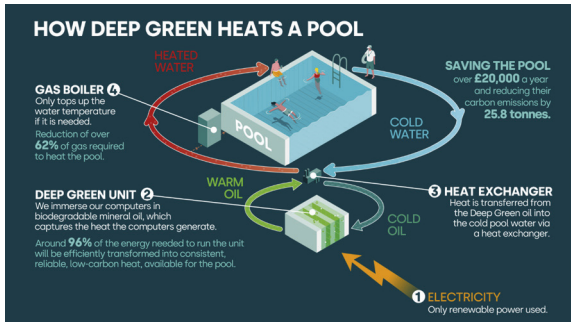


Typical layout of a water house building envelope, with water-filled glass.

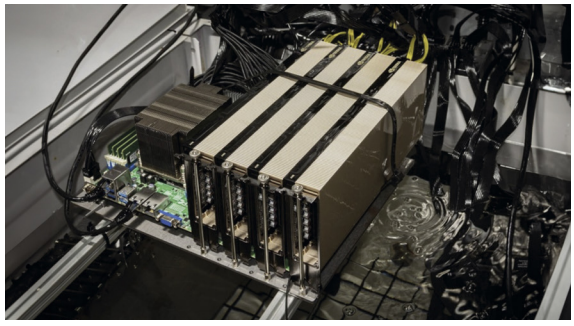
Retrofitting urban data centres with WFG panels and integrating algae cultivation could revolutionise their energy footprint, transforming them into 'regenerative data centers'. WFG panels could enable a single data centre to produce 400,000kWh of heat energy annually. Integrating algae systems further amplifies this impact, adding 80,693kWh of biofuel energy. Combined heat and power (CHP) plants are essential in converting this biomass into usable energy, making data centres not just energy-efficient but energy-regenerative.

Waste heat recovery is already being explored in colder regions like Finland and Sweden, where water is warmed to 50–60°C by the data centre cooling systems.[5] Gensler's research shows similarly that the average waste heat temperature in hyperscale data centres is 40°C. Crucially, this heat can be repurposed to produce water at approximately 30°C, a temperature perfect for cultivating *Chlorella vulgaris*, a microalgae species that thrives between 25–30°C.[6] One of the challenges of maintaining this biological fuel source is the requirement for consistent conditions to thrive.[7] The continuous 99.671% to 99.995% uptime of data centres ensures a reliable heat source for algae growth, seamlessly integrating biofuel production with existing infrastructure.[8]

In the UK, the Exmouth Leisure Centre serves as an example of innovative energy practices, where a digital boiler provided by the start up Deep Green offsets a significant portion of the centre's energy usage. The system uses a compact data centre, housed in a washing machine-sized unit, to generate heat from computing processes. This heat is transferred via oil to warm the pool, maintaining its temperature for 60% of the time and reducing energy costs.[9] Given the small scale of this operation, the example highlights the immense potential for much larger systems.



Deep Green diagram showing the heat exchange for the swimming pool.



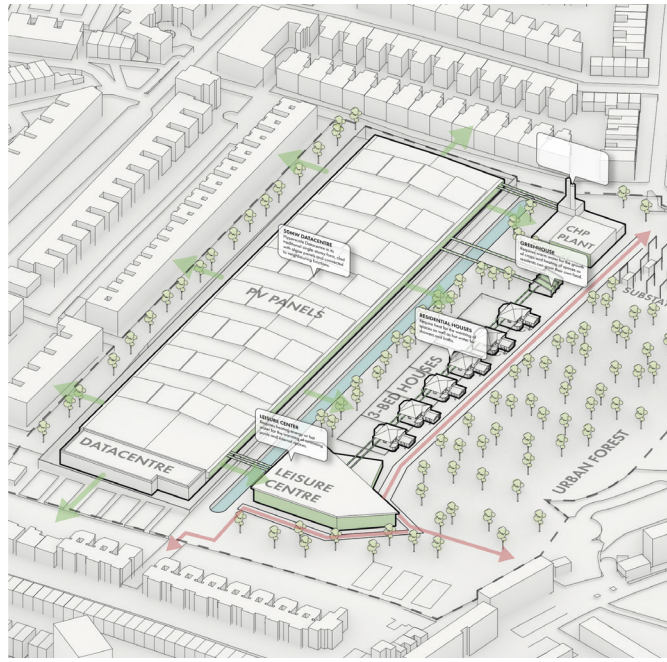
Deep Green server room or 'mini data centre'.

Another example of the potential of such systems is the Bio Intelligent Quotient (B.I.Q.) building in Hamburg, Germany, designed by Arup in 2013. Conceived as part of the International Building Exhibition in 2009, the project aimed to pioneer carbon-neutral and self-sustaining architecture for residential typologies. The B.I.Q. façade incorporates transparent panels filled with water and nutrients to grow algae year-round.[10]

This system offers a holistic synergy of solar energy capture, district heating and biomass production. Sunlight fosters algal growth, while excess heat is harvested through heat exchangers and integrated into a district heating network. Monthly, algae are extracted from the building's 129 photobioreactors, fermented and converted into biofuels intended for carbon-neutral electricity generation.[11]

While groundbreaking, the system exposed limitations, primarily its dependence on climatic conditions. Fluctuating light and temperature levels hindered efficiency, necessitating supplementary inputs that ultimately constrained scalability. Despite this, the example shows us that many of the constraints faced in residential buildings, particularly the constant low grade heat requirements, are more resolvable in data centres.

Scaling these solutions is becoming essential as the demand for data centres rises alongside increasing urbanisation. As the reliance on these buildings grows, so too does the potential of water-filled glass systems to exponentially enhance energy efficiency and production. Coupled with the continuous advancements in AI technology and operational efficiency, the future points towards a network of energy-generating networks capable of offsetting and even mitigating carbon emissions – ultimately transforming them into net-zero contributors to sustainability.



A data centre as an energy provider, with a long façade.

The long façades typical in these buildings presents the most strategic and impactful opportunity for an initial upgrade. This approach minimises disruption to server operations, which is a key concern for stakeholders, while leveraging the façade's extensive surface area and proximity to heat-generating elements. Additionally, given the sensitive nature of the proprietary server networks housed within, privacy is paramount. Privacy glazing solutions can also be specified in WFG modules. This innovative system not only enhances the thermal performance of façades – improving U-value and solar heat gain coefficient (SHGC) – but also provides dynamic control of light and visibility through tinted water.[12]

A pivotal vision for the future of energy production networks, powered by regenerative data centres, lies in avoiding default new-build construction that further contributes to immense carbon emissions and promoting retrofitting existing data centers and integrating onsite solutions within current structures. This strategy transforms existing data centres from passive energy consumers into active, sustainable regenerative hubs. The German government's mandate for all data centres to achieve a power usage efficiency (PUE) of 1.3 or lower has intensified the demand for retrofit solutions. This has spurred the development of advanced systems like direct-to-chip cooling.[13] However, relying completely on these more advanced technologies risks driving premature decommissioning of existing facilities to comply with the legislation. While such measures aim to enhance power efficiency and technological progress, the reliance on demolition and new construction could inadvertently lead to a substantial increase in overall carbon emissions, undermining the environmental benefits.



A data centre as an energy provider – residential tower close-up



A data centre as an energy provider - residential tower axonometric.

However, in some cases, existing solutions may fall short, making it imperative to design new builds that strike a delicate balance between energy demand and production. These developments must not only achieve net-zero emissions but, given the climate urgency, aim for net-positive impact, setting a bold standard for sustainable innovation. Take, for example, the creation of next-generation residential typologies seamlessly integrated with regenerative data server facilities. Such designs could not only generate the energy needed to power the homes they serve but also address urgent housing shortages.

The architecture, engineering and construction (AEC) sector plays a pivotal role in advancing sustainability. While it may not directly innovate in energy production technologies, the sector provides the physical and operational frameworks that enable these innovations to function effectively. Moreover, the industry has a significant opportunity – and responsibility – to rethink its processes through a proactive sustainability lens. Embedding decarbonisation strategies and rigorous life cycle assessments into the early stages of project planning can guide design decisions toward minimising carbon footprints and resource use. By prioritising net-zero or net-positive outcomes from inception, the AEC sector not only contributes to sustainable development but also establishes itself as a critical enabler of the global transition to low-carbon economies.

In this context, AI emerges as a vital tool for addressing sustainability challenges, albeit with important caveats. While it cannot independently solve systemic issues, AI offers unparalleled capabilities in processing vast datasets, identifying patterns and optimising complex systems. These strengths are particularly valuable in sustainability efforts, where data-driven insights are essential to measuring environmental impacts, predicting performance and refining processes. However, AI's utility depends on the quality of the data it is trained on and the clarity of intent behind its application. When applied thoughtfully, AI has the potential to improve decision-making, reduce waste and support environmentally responsible practices across the AEC sector in unprecedented ways.

This chapter has argued that, while AI is often criticised for its environmental costs, dismissing it risks overlooking its capacity to deliver profound environmental benefits. Scrutiny is essential – not to hinder its progress, but to guide its integration towards impactful outcomes. When rigorously applied, AI can transform environmental strategies, offering scalable precision and actionable insights previously unattainable.

With technological advancement inevitable, the urgency for ESG experts to lead AI integration is paramount. Their expertise is critical to ensure that AI's deployment aligns with strategic sustainability goals, balancing innovation with accountability. At this pivotal juncture, leveraging AI effectively is not merely an opportunity – it is a necessity to redefine environmental problem-solving at scale.

Notes

1. NVIDIA. 'Sustainability in Product Design and Recycling', https://www.nvidia.com/en-us/sustainability/, accessed 4 March 2025.
2. Forbes, 'The Key Role of Automation in AI Hardware Recycling', Forbes Tech Council, https://www.forbes.com, accessed 4 March 2025.
3. NVIDIA, 'Recycleye: AI Waste Recycling with NVIDIA Jetson', NVIDIA Blogs, (21 September 2021), https://blogs.nvidia.com, accessed 4 March 2025.
4. Z. Song, X. Zhang & C. Eriksson, 'Data Center Energy and Cost Saving Evaluation', *Energy Procedia*, 75 (2015), [doi: 10.1016/j.egypro.2015.07.178], pp. 1255–60.
5. M. Wahlroos et al., 'Future views on waste heat utilization – Case of data centers in Northern Europe', *Renewable and Sustainable Energy Reviews*, 82 (2018), [doi: 10.1016/j.rser.2017.10.058], pp. 1749–64.
6. J. Lansbergen & M. Redding, 'Using Algae to Offset the Carbon Cost of Data Centres', Gensler Blog (20 January 2023), https://www.gensler.com/blog/using-algae-to-offset-the-carbon-cost-of-data-centres, accessed 4 March 2025.
7. R. Serra-Maia et al., 'Influence of temperature on Chlorella vulgaris growth and mortality rates in a photobioreactor', *Algal Research*, 18 (2016), [doi: 10.1016/j.algal.2016.06.016], pp. 352–59.
8. Hewlett Packard Enterprise, 'What are Data Center Tiers? | Glossary', https://www.hpe.com/uk/en/what-is/data-center-tiers.html, accessed 4 March 2025.
9. Z. Kleinman, 'Tiny data centre used to heat public swimming pool', BBC News (14 March 2023), https://www.bbc.com/news/technology-64939558, accessed 4 March 2025/.
10. D. Perez, 'The First Algae-Powered Building Presents Unique Renewable Energy Solution', Engineering.com (11 December 2020), https://www.engineering.com/story/the-first-algae-powered-building-presents-unique-renewable-energy-solution, accessed 4 March 2025.
11. ARUP, 'Worldwide first façade system to cultivate micro-algae to generate heat and biomass as renewable energy sources', https://www.arup.com/projects/solar-leaf, accessed 4 March 2025.
12. M. Gutai & A. G. Kheybari, 'Energy consumption of hybrid smart water-filled glass (SWFG) building envelope' *Energy and Buildings*, 230 (2021), [doi: 10.1016/j.enbuild.2020.110508], p. 110508.
13. Dentons, Energy Efficiency Act: Relevance for Data Centers, https://www.dentons.com/en/insights/articles/2023/september/25/energy-efficiency-act-relevance-for-data-centers, accessed 20 December 2024.



Project Genesis: A Short Story on the Future Roles of Architects in 2040

Hamza Shaikh

DOI: 10.4324/9781003726500-19

It was 2040 and Mila's world had changed. As a design director in a leading global design firm, she worked with a team of brilliant minds – designers, technologists and sustainability experts – all working at the forefront of architecture's new evolution. They weren't just designing buildings, they were crafting experiences, ecosystems and solutions for a planet on the brink of renewed transformation.

Mila's day began in the firm's advanced design hub, a sprawling space filled with augmented reality interfaces, digital sketching tables, and screens that pulsed with generative models. Her team was deep into a project for a mixed-use tower that needed to meet ambitious climate and community goals. It wasn't just about aesthetics or function, every decision had to align with the firm's commitment to net-positive regenerative design.

She picked up her pencil and began sketching on paper, which she would transfer to the hybrid drawing interface – a tool that seamlessly pulled hand-drawn inputs into AI-driven expansion. Hayley, a young design intern, peering at Mila was surprised to see her still using old sketching methods and even more bewildered to see that Mila would draw with hard-pressing fast movements – 'No wonder her pencils are always blunt,' she thought to herself. Mila stood up and walked over to the virtual hub. 'Okay team, client update is in a few hours – what have we got?' she asked encouragingly. Kaylan was eager to show his design concepts first as usual, so he beamed his sketch into the virtual reality interface. Everyone activated their glasses to move into virtual space and were instantly transported into Kaylan's sketch world. He began at speed with digital hybrid sketches and paintings zooming around everywhere. 'Okay, check it – I've only got 273 ideas so far, but this one is epic. It's a naval-inspired avant-garde repositioning design where the tower, using this waste material from a nearby naval demolition project, helps to activate local retail regeneration for the community.' He continued, 'The familiarity of this historic naval building design would be quite a statement if we repurposed it for our tower. According to the community model simulation, people would love it – it got a green-65 score! Also, I just ran early carbon analysis and it shows a green-79 score, just for reusing the materials alone! This one's easy net-positive.' Arjun, the senior designer on the team, twiddling his digital pencil, looked sceptically at Kaylan's virtual avatar. 'What makes you think they would let us use this demo material?' 'Ah, I knew you'd ask that Arj,' replied Kaylan, confidently. 'Our site research bot caught this quick, and I validated it. Turns out they're the same contractors for our project, so I think we can put it forward as a valid upcycling proposition?' Arjun nodded slowly in approval, impressed with this information.

After only 30 minutes of deep concept exploration using sketch-to-virtual video experiences and agentic brainstorming, the team arrived at a solid concept. They collaboratively sketched out a new set of concept designs that were the product of thousands of ideas explored. 'Okay – Hayley, Lana, we are taking the multistorey food farming here,' Mila said, circling the new concept sketch in virtual space. 'And Kaylan, Arj, we've got the retail activation here for local communities, plus the solar façade all over here,' she said, finishing the rough but bold gestural sketch with smart labels. As she finished the final touches, her lines began morphing into more refined shapes. The software soon began approximating them into 3D forms. It analysed potential carbon emissions, daylight efficiency, wind flow and structural options in real time. It also linked directly to the local community memory model to give real-time data



insights and approval scores. A slight adjustment to the building's curve reduced its energy load by 22%, the software informed her. 'That was a good catch,' she murmured, as Arjun leaned over. 'Looks like the team is in agreement Mila – let's push options three and 19 into the materials hub to see if we can go fully carbon-negative on the cladding,' Arjun said. 'Yes, and let's forward the interior diffusion environment onto the clients ASAP. See if they want to play with the interiors with Kaylan for warm up – is he free?' Before Arjun could respond, a message notification popped up on their watches from Kaylan reading: 'Already done it LOL… Venv ready2go ;-).' Mila and Arjun laughed and joked that Kaylan, who the team nicknamed KAY.I, was secretly an AI super-agent.

The team worked like an orchestra, each member contributing their expertise working in multiple environments with hyperefficiency and collaboration. Lana, the sustainability lead, meticulously modelled the building's life cycle, ensuring it could be deconstructed and repurposed decades from now. Kaylan, who was their virtual simulation specialist, plus six other things, was refining the immersive VR environment that would allow the client to walk through the design as if it already existed. Every detail – from the texture of the flooring to the shadows cast by a neighbouring structure – had to feel authentic and backed by robust real-time data. There was no better team for the job.

By the afternoon, Mila joined the client presentation in the firm's immersive reality suite. The space transformed into the proposed tower's lobby, complete with sunlight streaming through its glass façade. The architects were joined virtually by the client and engineering team for the meeting and they moved through the digital 4D space together,

interacting with elements as if they were tangible. ‘This atrium doubles as a natural ventilation system,’ a voice explained, with areas being highlighted above showing the holographic airflow simulation. ‘It reduces cooling costs by 20% while creating a comfortable, open space for visitors.’ The clients were captivated by Eric’s data, the lead structural engineer giving input over the virtual call. But suddenly Jasper, a board member from the client’s team, interjected. ‘But hold on, please,’ he said, assertively. All virtual avatars turned their heads to the building lobby entry where Jasper’s avatar stood. ‘There’s something wrong here.’ There was a tense pause as every person in the room ran through all the possibilities of what might be wrong. ‘Didn’t we have the grand staircase more central here before?’ Mila replied, ‘Yes, well spotted Jasper – our systems showed us that by moving the staircase a few metres off-centre we were able to gain an extra 19% net internal area across all floors.’ ‘Really!? The staircase?’ Eric continued, ‘Well, the staircase moving led to a series of domino effects relating to the mezzanine layout and swimming pool. This led to amendments in the structural needs of the rear side of the tower.’ Eric ran a simulation for everyone to see what he was talking about as he continued to describe the series of events. ‘Of course, we wanted to limit as much steel use as possible, so we were also able to reduce the thickness of our structural columns.’ Mila chimed in with perfect timing: ‘As a result, we noticed right away that our rentable area already increased by 8.7%, so that prompted us to run another complete generative floor study optimisation, leading us to more than doubling our gains!’ There was a flurry of virtual thumbs ups and celebration confetti flying all over the virtual environment as the teams celebrated a very successful meeting.

The following week, Mila gathered her team to discuss an issue that was flagged in the generative BIM standards feedback report. ‘Okay all, we’ve seen the complaint we received from a senior member this morning – what’s going on?’ They think it’s a dataset issue with our i23D model,’ said Mila with concern on her face. Farid, the technical data engineer, adjusted his glasses and cleared his throat before replying nervously: ‘Yes – I think they are correct… I’ve looked deeper into it and our Approxi-M8 model has not dealt with artefacts and hallucinations from the reconstructed model properly…’



He cleared his throat again and continued, 'It appears to be converting tiny undetected artefacts into random object assets – like children's toys.' Kevin with his usual humorous timing abruptly and loudly jumped in, 'THAT'S why I kept seeing spoons and rubber ducks in the model!' There were quiet giggles around the table. Mila looked over at Kevin and he immediately stopped smiling. 'Kev, why did you not mention this?' 'I'm sorry, I thought it was an issue from my imports, so I cleaned up all the weird objects and thought nothing of it,' replied Kevin. Farid continued, "It's a synthetic feedback loop creating hallucinatory assets caused by recursive generative degradation between the 2.5D and 3D LLM conversion algorithm…' There was silence. 'Yes, that was my other theory…' said Kevin again, causing more laughter. 'Essentially, I need to fix the script to cull artefacts at a lower tolerance. It's good we caught it now – I can try to fix it up by lunch time,' concluded Farid. Mila sighed with some relief and concluded the meeting by reminding the team that hybrid approaches were paramount, ensuring human intuition, oversight, communication and input were embedded at every level.

Arriving to the Generative Construction Site

The project work continued and a few months later, Mila arrived onsite with her team to witness the next phase of their vision coming to life. The atmosphere buzzed with activity as they stepped out of their electrified vehicle into the sprawling construction site. Towering robotic arms moved with uncanny precision, assembling the building's modular components from onsite generative systems. Silent drones hovered overhead, capturing real-time data and monitoring every movement. Every step was choreographed by the digital twin system, a virtual replica of the project that tracked progress and flagged potential issues before they could arise. Their smart helmets gave them real-time audio updates through computer vision integrations. Metallic clinks and humming mechanical movements could also be heard as the building was in pre-assembly stage, but the site was quiet enough to hear some distant footsteps approaching.

'Hey, the dream team is here – welcome!' exclaimed the site manager. A seasoned contractor named Samir greeted them and led them toward the central monitoring

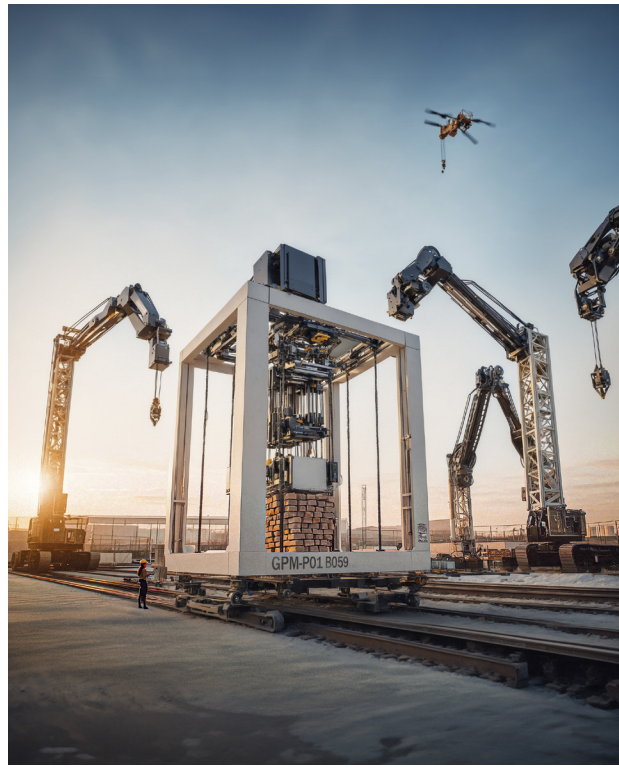


station. Screens displayed a wealth of real-time data; build progress was on track at 12% completion, carbon emissions had been further reduced by 21% compared to initial projections and procurement updates confirmed that all materials for the next phase were already onsite. Mila looked down at the digital timeline on her watch – they were three weeks ahead of schedule thanks to the recent updates the team had made on their 3D-to-build software, made in collaboration with Samir's team. Samir was also looking at his watch. 'The 3D-to-build update is fantastic! Your team made huge inference efficiency gains… we were impressed.' He continued, 'But now let me impress you! Come this way.' His team led the group to the systems hub where a glowing table of data infographics beckoned. Everyone in the room rushed over and began pulling and pushing the holographic NeRF projections to explore the design and data. Hayley the junior intern exclaimed, 'Woah, the core design enables a 98% egress efficiency! I didn't even know it could go that high!?' 'Neither did we,' said Samir. 'But the extended procedural planning exercise we did was a huge help, and once we integrated it with the digi-twin behavioural simulation, we cracked it.' Samir put his finger to his ear pushing a button on his earbud, 'Genesis. Run digi-twin simulation replay.' The holographic projection suddenly flickered to a 3D animation of the floor simulation studies with little people walking around the model – the team were stunned.

Suddenly, there was an extended beeping on Samir's watch and a synchronised red light flashing on the system hub ceiling. He looked to see a health risk alert flashing across the displays. It was flagging an incoming extreme weather event, forecasted to hit the area in 26 days' time. 'Sorry team, we have another storm warning. Looks like it's starting in less than a month – we gotta go. Please, do head down – we will catch up with you in an hour or so.' On the way out, Hayley walked deliberately slowly to overhear the team's strategy. 'Okay, response options conclude that we'll need to secure the site and accelerate key structural tasks,' Samir explained. His team member updated the schedule on the fly using the holographic display. 'I am redistributing tasks to ensure critical milestones are met before the storm's arrival,' said his colleague. The AI-assisted planning system worked seamlessly in parallel recalibrating the workload, coordinating shifts and resources to maintain efficiency while minimising risk. 'Super-cool,' murmured Hayley with a grin as she made her way downstairs.



With the revised schedule underway, Mila and her team moved on to observe the robots in action. One of the large robotic arms meticulously applied precision joints to prefabricated panels, while another handled the integration of solar components directly into façade panels. But the most exciting robot was the enormous square-framed generative parts and materials robot which was printing and stacking parametric biobricks and bioblocks with eerie speed. Each brick was different and perfectly stacked in coded areas along the edge of the site. As the bot moved smoothly on a track around the site, Hayley stared in awe and asked Mila, 'They made this thanks to your research, right? I've studied it!.' She stood mesmerised by its speed and lack of noise. 'It can print almost any bio-industrial material in any shape or form required... Plus all systems mitigated 84% of transportation emissions and 75% regional mining volume this year so far,' said Hayley, showing off her knowledge. Mila, who was smiling in pride as she captured footage on her glasses continued, her voice tinged with excitement, 'Thanks to generative algorithms and material science breakthroughs, it doesn't just print, it optimises. Every biomaterial is uniquely designed for its load-bearing role and environmental conditions, using locally sourced waste and carbon-captured composites. This gives us a huge boost in getting to net-positive because all carbon emissions that we can't remove go into sequestration and bio-material production for carbon-fibre reinforcement – we had five new carbon capture contracts this year.' 'Oh, and,' she added with a smirk, 'this is just the prototype.' Mila smiled and continued capturing footage while Hayley stared at her in awe.



Before leaving, Samir informed Mila that despite the storm response plan, work had to stop for six days and he would update the wider team later. Mila thanked Samir and the team, then stood for a moment at the edge of the site, watching the symphony of machines and people working in harmony. She paused and reflected for a moment, feeling light raindrops and wind gusts brush across face. She humbly thought to herself that, even in the age of generative AI and robotics revolution, the earth and nature dictated our true future. As she pondered all that had changed over the years, the reality of the project from the early sketch brainstorm sessions and carbon analysis to the profound electrified hums of robotic activity, all within a space of a few months, it gave her a sense of accomplishment. This was the architecture of 2040: hybrid, adaptive and technology deeply rooted in environmental purpose.

As they were ready to leave, the weather warning continued to weigh heavily on Mila. It was the third storm in the last few months, but Mila was optimistic that their firm's diligent collective efforts would begin paying off soon – from their research lab initiatives to the thousands of projects they touched every year globally. She was reminded that the resilience team's predictive model had forecasted that this autumn storms were going to be frequent, and that the firm had established their groundbreaking regenerative systems 15 years ago, ensuring their designs didn't just minimise harm but actively contributed to positive ecological restoration. She felt a little more at ease when she thought of the new tower's multistorey roof farm hosting urban agriculture for the residents, and its façade's solar-integrated panels that would generate surplus energy for the local grid. Just then, Hayley slowly approached Mila and with a timid smile said, 'Hey – hi… me again.' Hayley laughed nervously. 'Just wanted to say, I really look up to you Mila – I've been following your work for many years; I even did my dissertation on you! Well, not you personally. Actually, well, kinda… But more on your research! On the GPM bot and 2D-to-3D data… you know.' She laughed nervously again. Mila smiled gently and responded, 'That's so sweet of you to say, Hayley. My hunches are pretty good and I feel you have a very bright future. Hopefully it's with us!' Hayley's nervous smile turned into a more comfortable and confident one. But she then looked immediately thoughtful and asked, 'Okay weird question: Why do you still sketch on paper? And why use a blunt pencil?' Mila looked jovially at Hayley and, laughing, she said, 'I'm surprised you noticed I prefer blunt pencils! The truth is Hayley, nothing beats your hand holding a pencil. It's the most truthful way to portray your ideas, it's human – and we can't lose that…' 'Right!' approved Hayley, 'but why a blunt pencil?' Mila began to turn to catch up with the group who were leaving the site. She gently tapped Hayley's head and said: 'Because, we need bold ideas.'

As the work day came to an end, Arjun and Mila sat on the wellbeing terrace together following wind-down time. 'What a day,' said Arjun. 'So much has changed since we were training!' Mila nodded calmly. 'Do you remember what our sustainability tutor told us? That every decision we make as designers is part of a 'Grand Permanent System' – GPS she called it, our navigation system,' he laughed. 'I never believed her… But I do now.' Arjun looked over the terrace view onto the city in the distance. 'I guess I never felt part of the bigger picture before. But now I can see how everything we do, the materials we visualise with, the forms we draw, the words we use, the buttons we press – every decision makes a difference.' Arjun paused and looked at Mila. 'I thought it was childish to think you can change the world. I wish I'd believed her…' Mila looked over to Arjun and quietly said, 'Me too.'

As the day wound down, Mila returned home feeling a sense of deep responsibility and purpose. She reflected on the project's journey from sketch to simulation, from algorithms to robotic assembly. Suddenly an electronically friendly voice interrupted with a polite notification noise: 'New E-mail from EkEye systems – it's a revised job offer for Chief Hybrid AI Officer leading human integrations for their Artificial General Intelligence design model. They have doubled their salary offer and tripled bonuses. Shall I delete this one too Mila?' asked Mila's assistant-bot. Mila paused for a moment. She then walked slowly over to the kitchenette, picked up her herbal tea that was pre-prepared and at the perfect temperature. 'Start bedtime vibes,' she instructed the bot. The lights dimmed to an ember orange hue, lavender aromas filled the room and furniture adapted.

As Mila sipped her tea and listened to the surrounding sounds of Tchaikovsky's *Sleeping Beauty*, Op.66, Act I 'The Spell', she began roughly sketching a new concept for the Project Genesis – National Space Station, commencing next week. Her apartment windows had activated a filtered tint to darken all city light wavelengths, letting in all the starlight, and a smart silent frequency to silence all city noise. The vibrant music continued in the background, helping her mind to stop racing as she playfully sketched remarkable and extraordinary ideas. Her bold pencil marks were in unison with the broad violin strokes. Holding her blunt pencil, she gazed up at the stars and felt the limitless potential they represented.



ABOUT THE CONTRIBUTORS

Aleksandar Bajic is a RIBA Silver Medal and AJ Sustainability Award-nominated architectural designer, recognised for his expertise in sustainable innovation and carbon-neutral solutions. As part of Gensler's Building Transformation & Adaptive Reuse studio, he applies forward-thinking strategies that reshape the built environment, focusing on resilience and the thoughtful integration of sustainable practices. With international experience in visionary design, Aleksandar plays an integral part in transformative reuse projects, spearheading innovative strategies that redefine the boundaries of architectural design and promote both environmental responsibility and forward-thinking solutions.

Carlos Bañón is Associate Professor of Architecture at the Singapore University of Technology and Design (SUTD) and Director of its Architectural Intelligence Research Lab (AIRLAB), where he leads pioneering work in AI-driven structural design, advanced digital fabrication and geometric sustainability. He co-authored the influential *3D Printing Architecture: Workflows, Applications, and Trends* and consistently bridges cutting-edge computational technology with human-centric architectural practice. His forward-thinking designs have earned him Singapore's President's Design Award and a nomination for the prestigious EUMIESAward. His projects have been showcased at the Venice Biennale of Architecture and honored with the Design of the Year prize by the London Design Museum, underscoring his global impact as a thought leader. As an advocate for AI as a creative collaborator, he champions a design ethos where technology augments human ingenuity rather than replacing it.

Matias del Campo is an architect, designer and theorist, currently serving as Director of the MS ACT program (Architecture, Computational Technology) and as Associate Professor at the New York Institute of Technology (NYIT). He co-founded SPAN in Vienna in 2003 with Sandra Manninger, establishing a practice renowned for its integration of contemporary technologies in architectural production. SPAN's award-winning designs are shaped by the intersection of computational methodologies and philosophical interrogations, a conceptual framework they describe as 'design ecology'. Matias del Campo's innovative contributions have been recognised through prestigious awards, including the Accelerate@ CERN fellowship, the AIA Studio Prize and the ACADIA Innovative Research Award of Excellence. His work is included in the permanent collections of notable institutions such as the FRAC, the MAK in Vienna, the Benetton Collection, the Pinakothek Munich and the Albertina. He has authored several books on AI and Architecture such as *Neural Architecture* (ORO), *Diffusions* (Wiley), *Machine Hallucinations* (Wiley) and *Artificial Intelligence and Architecture* (Wiley).

Sara Codarin is an architect and educator at Lawrence Technological University's College of Architecture and Design. Her research spans robotic fabrication, large-scale additive manufacturing, digital craft, computational workflows and generative AI for world-building. She holds an M.Arch and a Ph.D. from the University of Ferrara, where she focused on robotics for cultural heritage preservation. Her collaborative work on AI was selected for the 2025 Venice Biennale of Architecture, curated by Carlo Ratti. She has exhibited her work at the Pratt Institute's School of Architecture, Detroit Month of Design and Milan Design Week. Sara is a techno-optimist and she loves robots.

Rada Daleva is an architectural designer and entrepreneur. At Studio Tim Fu, she leads groundbreaking architectural projects that integrate AI into the design process. At this innovative London-based architectural practice, she spearheads a multidisciplinary team that leverages advanced AI and computational design to transform complex ideas into sustainable, visually striking built environments. As the founder of Daleva Design, she develops high-end biomorphic furniture that reflects the intersection of nature, technology and craftsmanship. With a background in architecture and a fascination with AI and biophilic design, Rada is driven to explore ways to combine these fields to create better living environments where humans can thrive.

Tim Fu is a renowned designer known for his use of advanced computation and AI. Emerging from Zaha Hadid Architects, he has since founded Studio Tim Fu, high-tech architectural practice pioneering in the integration of AI into visionary designs. As a prominent voice in the field, his work has been showcased at conferences and exhibitions worldwide. Leading a specialised team of architects and technologists, he undertakes projects around the globe. Through his design and technological insights, Tim has amassed a significant online following, establishing his influence as a thought leader in the industry.

Will Garner is a designer best known as an early adopter of AI in design, emphasising its transformative power as a creative design collaborator. Starting with initial experiments on early AI diffusion models, Will has built a significant social media presence showcasing the use of AI in design. Currently, he leads a design studio focused on empowering individuals and organisations to harness AI in their design processes. Over the past few years, Will has applied AI to a broad range of projects, from large master plans to live production visuals.

Ömer Nuray is an architect and founder of Design Input Studio, where he develops practical resources to help architects integrate technologies like AI and computational design into their workflows. His work focuses on making advanced design tools accessible and use ful for the architectural community. In his role as an AI Architect at Rendair AI, Ömer is involved in building AI-driven visualisation tools that simplify the process of creating high-quality architectural visuals. Through Design Input Academy, his YouTube channel, and social media, Ömer shares his knowledge and experiences, offering architects and designers clear and practical guidance on using AI in the design process

Hassan Ragab is an interdisciplinary artist, architect and computational designer whose practice explores the intersection of artificial intelligence, architecture and visual culture. His work critically investigates themes of identity, heritage and the transformative role technology plays in shaping creativity, aesthetics and perception. Drawing from over 15 years of experience in architecture and computational design, Hassan integrates generative AI tools to challenge conventional visual narratives and provoke dialogue on contemporary societal and cultural issues. His projects often blend digital art, research and storytelling to inspire deeper reflections on human values and collective identity.

Keir Regan-Alexander is an architect and AEC domain expert operating with one foot in practice and one in live software development. Following a directorship at London's Morris+Company, Keir founded Arka Works to prepare the architectural profession for AI-driven transformation. He specialises in integrating generative AI tools into professional workflows, helping businesses adapt to a rapidly evolving digital landscape. In 2024, he co-founded OmniChat.uk, an enterprise-grade, multimodal LLM platform which is now deployed in practices across Europe.

Ismail Seleit is a design architect at Foster + Partners, specialising in design technology. With expertise in architecture and computational design, he has contributed to high-profile projects spanning competitions to large-scale built environments. He plays a key role in advancing design workflows, leveraging BIM and computational tools to drive efficiency, innovation and cross-disciplinary collaboration from urban planning to product design. His research explores the frontier of generative AI, particularly Stable Diffusion and ControlNet, to pioneer controlled image composition techniques. Beyond architecture, Ismail is an accomplished ambient-electronic music producer, composing film scores for architectural shorts and independent films. His passion for innovation and creative problem-solving continues to shape his multidisciplinary pursuits.

Thomas S. Sibley is a London-based electronic artist and inventor who designs and develops spatial projection technology for hybrid reality environments. His latest installation, *Alien Antenna*, premiered in February 2025 at AADK's Centro Negra in Blanca, Spain and applied new forms of simulated depth to 360-degree photographs, digital collage and geometric forms. He is designing a spatial theatre system that projects volumetric simulations upon real-world environments.

Kat Stevens is a RIBA Part 3-qualified architect based in Dubai. She previously worked at Zaha Hadid Architects in London and studied parametric design at the University of Westminster. She has a deep interest in biophilia, wellness and human-centric design, and explores how AI can aid design processes beyond generating images. Her experiments have led to new ways of integrating AI into design, some of which are featured in this book. Now in Dubai, she continues to explore the integration of AI, nature and evolving technologies in architecture.

INDEX

#

3D modelling
  traditional **83, 124, 125**
  see also generative 3D modelling
3D printing **6**
3D visualisation **5**
  extended reality (XR) **168, 175**
  linear perspective **168-169**
  ray tracing **34**
  stereoscopy **173-174**
  virtual reality **5, 168, 174, 175**
4D visualisation **170-173**

A

Abbott, Edwin Abbott **170, 171**
accessibility **5**
ADE20K colour scheme **127, 129**
Adobe **197**
agentic systems **X**
Agrawala, Maneesh **127**
AI-driven architectural workflows **124-135**
  case studies **129-132**
  layers of **127-128**
  versus traditional architectural workflows **124-125**
'AI Form-Morphing' workflow **109-120**
'AI Stone Column' project **86, 87**
algae cultivation **202-205**
AnimateDiff **63**
Appel, Arthur **34**
architects *see* role of the architect
Architects' Journal **107**
architectural practice
  beginning of AI revolution in **80-87**
  effects of AI on **102-107**
  organisational restructuring **87**
  reduction of workforce **87**
  see also AI-driven architectural workflows
artificial general intelligence (AGI) X, **102-103**
augmented reality **5**
Autodesk **196, 197**
Automatic1111 **XIII, 62, 110-119**
automation **6-7**
autostereoscopy **168, 172, 175**

B

Bezos, Jeff **103**
Bio Intelligent Quotient building, Hamburg, Germany **203**
biofuels **202-205**
biomorphic furniture **67-76**
Blender **10-15**
'Breeding New Typologies' workflow **89-98**
Brunelleschi, Filippo **168**
building information models (BIM) **84, 107, 133, 192, 193-194, 196**

C

ChatGPT **X, 92, 138-147, 180-187**
circular economy **200**
Claude 3.5 **102**
climate challenge **200-206**
Collab notebooks **62**
collaboration **5, 196, 197**
combined heat and power (CHP) **202**
ComfyUI **XIII, 63, 104**
  workflows **10-11, 42-53, 95-96, 138-147**
computer-aided design (CAD) **59, 107**
  see also Rhino
computer-vision approximation **X**
construction industry **5-6, 7, 197, 213-216**
ControlNet **XIII, 5, 42, 82, 127-128, 133, 134**
  workflows **15, 42, 45, 95, 140-143**
customisation **105-106**
cyclical energy production **200-201**

D

DALL·E **X, 3-4**
data centres, environmental impacts **200-206**
data-driven AI systems **192-197**
deconstructivism **36-38**
deep dreaming **34, 37**
Deep Green **202**
Democritus **32**
depth maps **127, 129, 130, 131**
Derrida, Jacques **36-38**
Differential **194**
differentiated localised materiality **131**
diffusion models **34-35, 36, 37, 38, 58, 83-84, 124, 133**
  see also Midjourney; Stable Diffusion
diffusion of innovation theory **107**
digital twins **5, 213**
direct-to-chip cooling **204**
district heating **203**
Dürer, Albrecht **34**

E

echolocation **32**
efficiency **2, 6-7, 83-84, 193, 194, 196**
efficiency simulations **133**
Einstein, Albert **171**
Eisenman, Peter **36, 38**
energy demands of AI **200-206**
environmental impacts **200-206**
Euclid **32**
Exmouth Leisure Centre, UK **202**
extended reality (XR) **168, 175**
extramission theory of perception **32, 36**

F

fiction, representation as **36, 38**
file formats **196**
Finch3D **194, 195**
Flatland: A Romance of Many Dimensions (Abbott) **170, 171**
Flux **3-4**
FormFinder model **110-119**
four dimensions, visualising **170-173**
Fourth Dimension, The (Hinton) **170**
furniture design **67-76**

G

game theory **105**
Gaussian splatting models **10**
'Generative 2D to 3D' workflow **9-16**
generative 3D modelling **7, 105, 192-195, 197**
  workflows **10-15, 42-53, 71-74, 129, 152-163**
generative 3D printing **6**
generative AI technology, evolution of **X-XV, 6-7, 59-65**
generative design **103-106**
generative pre-trained transformer (GPT) **102, 176**
  see also ChatGPT
'Generative Synesthesia' **103**

Giotto 168
Goldstein, Robert A. 34
GPT-3 102
GPT-4 176
Grasshopper 26, 60, 104, 197
  workflows 42–53, 180–187

H
Hadid, Zaha 80
hallucination, in representation 36, 37–38
Harari, Yuval Noah 65
Heidegger, Martin 37, 38
higher dimensions, stereoscopy 173–174
higher dimensions, visualising 168–176
Hinton, Charles Howard 170
holograms 168
Huang, Jensen 65
hybrid AI approach XIV–XV, 4
  see also AI-driven architectural workflows
hypercubes 170, 172–173
hyperspatial architecture 176
hyperspatial theater 168, 175

I J
'I Prompted a Sofa, then Sold It!' workflow 67–76
image prompt adapters (IP -Adapters) 127–128, 131, 132
ImageNet dataset 34
immersive technologies
see extended reality (XR); virtual reality
inclusivity 5
industrialisation 59
Industry Foundation Classes (IFC) 196
Information Age 174–175
inpainting 128, 130, 132
intromission theory of perception 32

K
Kohya 94
Krita 152–163
KSampler 45–46, 141–142

L
large language models (LLMs) 38, 102, 105, 107, 110, 180, 192
  see also ChatGPT
large parameter models 102
large reconstruction models (LRMs) 10–15
'Laser Cut Prompts' workflow 19–28
life cycle assessments (LCAs) 200, 206
linear perspective 2, 168–169
LookX 82, 83
low-rank adaptation (LoRA) training 90–97, 105, 110–119, 133, 134, 152

M
machine learning 6, 10, 34, 35
McNeel 197
MDF laser cutting 20–28
Midjourney X–XII, 3–4, 60, 61, 62, 63, 81, 102
  workflows 20–27, 68–75, 90–97
mimesis 33, 36
'model collapse' 90
modernism 59
Mostaque, Emad 65
multi-layered reality perception 176

N
Nagel, Roger 34
net-positive 200, 206
neural networks X, XIII, 127, 176
neural radiance fields (NeRFs) 10, 193
NVIDIA X, 5, 197, 200

O
Objaverse dataset 10
OpenAI X, 64, 102
  see also ChatGPT
outpainting 128

P Q
parametric techniques 103–104
Parc de La Villete 36
'peak novelty' 103
perception 32–35, 36, 38
personalised AI models 133, 134
perspective, linear 2, 168–169
Petals model 110–119
photogrammetry 10
physical modelling 124, 131
Plato 32
platonic solids 170–171
PNAS Nexus 103
point cloud meshes 193
polytopes 170–171
poststructuralism 36
predictive deep learning 84
premortem technique 104–105
Processing 4.0 138–147
Project Genesis 209–217
'prompt-crafting' 81
Ptolemy 32
Python 180–182

R
Rao, Anyi 127
rasterisation 34, 38
rasterised 3D assets 192, 193
ray tracing 34, 38, 84
'Real-Time AI-Enhanced Digital Collaging-to-3D' workflow 151–164
'Real-Time Façade Iterations' workflow 41–54
real-time visualisation with AI 42
reality
  augmented 5
  extended 168, 175
  multi-layered reality
perception 176
  representation and 36–38
  virtual 5, 168, 174, 175
recursive training 90
Renaissance 2, 38, 170
Rendair 82
rendering 34, 36, 38, 84, 124
representation 33–37
resource demands of AI 200–206
Revit 194
Rhino 60, 129, 194, 196–197
  workflows 10–15, 20–27, 68–75, 152–163
  see also Grasshopper
risk reduction 5, 6–7
'Robotic Mark-Making' workflow 179–188
robotics-enabled construction 5, 6, 197, 213–216
Rogers, E. M. 107
role of the architect 5
  history of 58–59
  impacts of AI on 7, 59–65, 106–107, 125, 133
  story on future of 209–217
  visions of future of 17, 29, 55, 77, 99, 121, 149, 165, 189
Roth, Scott 34
Royal Society 173

S
'Sculpting with Inference' workflow 137–148
SDXL-Lighting model 42

INDEX

SDXL-Turbo model **42**
segmentation **127, 128, 129, 138, 140–141, 142**
semiconductor industry **200**
Seth, Anil **36**
Shoes (van Gogh) **36, 37**
sofa design *67–76*
solar energy capture **203**
Sora **64**
spatial projection technology **175**
Speckle **194–195**
Stability AI **10, 65**
Stable Diffusion **XII–XIII, 3–4, 61, 62, 63, 82, 102, 104**
  workflows **10–11, 15, 42–53, 95–96, 110–119, 138–147, 152–163**
  see also ComfyUI; ControlNet
stereoscopy **173–174**
Stott, Alicia Boole **170**
structural analysis **133**
style transfer **37**
sustainability challenges **200–206**

T

Tesseract Project **168, 175**
tesseracts **170, 172–173**
text-to-image generation **34**
  see also Midjourney
textual prompts **127**
TripoSR **10–15, 152–163**
Truth in Painting, The (Derrida) **37**

U

Unity **5**
universal file formats **196**
Unreal Engine **5, 193**
urban low-energy networks (ULEN) **201**

V

van Gogh, Vincent **36, 37, 38**
vectorised assets **192, 194**
videogrammetry **10**
virtual reality **5, 168, 174, 175**
visual culture, history of **2–3**
Vredeman de Vries, Hans **168**

W

WarpFusion **63**
waste heat recovery **201–205**
water-filled glass (WFG) technology **201–205**
Wheatstone, Charles **173–174**
Whitted, Turner **34**
workflows
  'AI Form-Morphing' **109–120**
  'Breeding New Typologies' **89–98**
  'Generative 2D to 3D' **9–16**
  'I Prompted a Sofa, then Sold It!' **67–76**
  'Laser Cut Prompts' **19–28**
  'Real-Time AI-Enhanced Digital Collaging-to-3D' **151–164**
  'Real-Time Façade Iterations' **41–54**
  'Robotic Mark-Making' **179–188**
  'Sculpting with Inference' **137–148**
  see also AI-driven architectural workflows

X Y

XR *see* extended reality (XR)

Young, Michael **36**

Z

Z-depth maps **127, 129, 130, 131**
Zaha Hadid Architects **80**
Zhang, Lvmin **127**

IMAGE CREDITS

Hamza Shaikh
viii, xi–xvi, 3–4, 6, 8, 11–15, 18, 108, 150, 153–163, 190, 197, 199;

Kat Stevens
21, 22–27;

Matias del Campo
30;

Johannes Zahn, 1687
32;

SPAN (Matias del Campo, Sandra Manninger), 2023
33;

Ömer Nuray
40, 43–53;

Hassan Ragab
56, 58–65;

Rada Daleva
66, 69–70, 72–75;

Tim Fu
78, 81–83, 85, 87;

Tim Fu and Till Apfel
86;

Photograph by Graham Uden/Fortune
80;

Will Garner
91–97;

Keir Regan-Alexander
100, 104, 106;

Ismail Seleit
111–119;

Carlos Bañón
122, 124, 126, 128–132, 134, 136, 139–147;

Tulsi Bhagat, CC BY 4.0
166, 170;

Thomas S. Sibley
168, 172, 175;

Universitäts- und Landesbibliothek Düsseldorf, CC BY 1.0
169;

David Eric Ffell, CC BY 4.0
171 (top);

Public Domain
171 (bottom), 174;

Wolfram Research, Inc, CC BY 3.0
173 (top);

Adapted from Nityananda & Read (2017), *Journal of Experimental Biology*, CC BY 3.0
173 (bottom);

Sara Codarin
178, 181–187;

Speckle
195 (top);

Finch3D
195 (bottom);

Water-filled Glass Ltd
201;

Deep Green
202;

Aleksander Bajic
203–205.

© RIBA Publishing, 2025

Published by RIBA Publishing
66 Portland Place
London W1B 1AD

ISBN 9781915722393
DOI: 10.4324/9781003726500

The rights of Hamza Shaikh to be identified as the Editor of this Work have been asserted in accordance with the Copyright, Designs and Patents Act 1988 sections 77 and 78.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior permission of the copyright owner.

British Library Cataloguing-in-Publication Data.
A catalogue record for this book is available from the British Library.

Commissioning Editor: Alex White
Project Editor: Anna Watson
Assistant Editor: Patricia Roig
Production: Marie Doinne
Designed and Typeset by CHK Design
Printed and bound by Pureprint Group Ltd
Cover image: Hamza Shaikh

While every effort has been made to check the accuracy and quality of the information given in this publication, neither the Author nor the Publisher accept any responsibility for the subsequent use of this information, for any errors or omissions that it may contain, or for any misunderstandings arising from it.

No part of this book may be used or reproduced in any manner for the purpose of training artificial intelligence technologies or systems.

For product safety concerns and information, please contact our EU representative:
Taylor & Francis Verlag GmbH
Kaufingerstraße 24
80331 Munich
Germany

Email: GPSR@taylorandfrancis.com

www.ribapublishing.com