Artificial Intelligence, Perception and Creativity

How deep neural networks can be used to augment human agency in the design process.

Tutors: Carolina Ramirez-Figueroa, Luis Hernan
The Bartlett School of Architecture, UCL
Design Thesis Report
29.07.2021

Chris Whiteside

Abstract

This thesis will explore the potential of deep neural networks to challenge human perception and augment the creative process of architectural design. It will begin by analysing the potential of Artificial Intelligence (AI), comparing machine intelligence, perception and creativity to the human counterparts on which these terms are based. This comparison will help to identify the strengths and weaknesses of both human, and artificial, intelligences and how they can supplement each other to create an augmented design process. Through case studies of pioneering neural architects, it will demonstrate how neural networks can build their own perception and challenge the human perspective. These examples will be used to inform my own experiments with Pix2Pix GANs, which will assess multiple methods of designing and curating training datasets, and how these can be used to build design tools powered by deep learning.

Keywords: Artificial Intelligence, Deep Learning, Neural Networks, Creativity, Perception, Hallucination

Image: Early GAN Synthesised Image. Author.

Table of Contents

Introduction

Artificial Intelligence (AI) arose as an attempt to recreate intelligent entities from the natural world, such as the human brain, and has captivated popular interest since its early representation in science fiction (Russell and Norvig 2010). Throughout its history, AI has been pitted against our own human intelligence, from the Turing Test to world title chess matches, and dystopian predictions of machine takeover. Whilst early AI aimed to replicate the human mind, today the focus has moved to Narrow AI, which can excel to super-human levels of intelligence in specific tasks, but lacks the general intelligence of humans. AI remains a technological tool like the human inventions which came before it, and its potential lies not in replacing human intelligence, but in augmenting it.

It is deep learning, powered by artificial neural networks, that offers the greatest potential and has powered the AI revolution in many industries, from e-commerce to autonomous vehicles. This has resulted in huge increases in productivity through automated management and distribution systems, yet is also spawning new challenges surrounding data privacy and bias. In the field of architecture, these effects have not yet been felt, with the application of deep learning limited to a few pioneering firms and academics. Of particular interest to architectural designers are Generative Adversarial Networks (GANs), which transform AI from a purely analytical tool to a generative agent (Chaillou 2019). By a process of pattern recognition, neural networks are able to build their own perception. They are able to view and analyse data in a way that can be compared to our own human perception. Our perception contains preferences and prejudices built up through our past experiences. Can the machine perception of neural networks be used to challenge our own preconceptions? Can it augment and extend the creative potential of architectural design?

This thesis will analyse machine intelligence, perception and creativity in order to form a response to these questions. It will use case studies of pioneering neural architects, as well as a series of design experiments exploring the application of neural networks. In doing so, it will investigate various methods of curating, designing and augmenting datasets for use with deep learning, and will speculate on the future potential of neural networks within architectural design.

Russell, S., Peter N. 2010. Artificial Intelligence: A Modern Approach. 3rd ed. Prentice Hall. | Chaillou, S. 2019. 'The Advent of Architectural AI'. Harvard Graduate School of Design.

Intelligence, Perception and Creativity

Intelligence

Intelligence - is this an inherently human concept? Superior intelligence has long been considered the characteristic of the human species that sets it apart from the rest of the living world. It was in 1758 that Carolus Linnaeus classified humans as *Homo sapiens* or in Latin “wise man” (Tattersall 2020). Human intelligence is what led to the creation of the first stone tools, to the domestication of plants and animals, and the settlement of the human race. It is what led to technologies such as the steam engine, powered flight, wireless communication and the internet. But where does this intelligence come from? How do we learn? How were these ideas generated? These are the questions many are asking today, with one main goal in mind - the creation of intelligent machines.

The field of AI arose as an attempt to understand and recreate intelligent entities from the natural world, such as the human brain (Russell and Norvig 2010). First coined in 1955, AI has been evident in popular culture for much longer, from the robotic helpers in the *The Iliad* of ancient Greece, to the Tin man in the *Wizard of Oz* (Akst 2020). By the time Alan Turing began to explore the mathematical possibility of intelligent machines, they were already a staple of science fiction. His seminal paper, Computing Machinery and Intelligence, published in 1950, addressed the question, “Can machines think?” and considered how to test machine intelligence against our own, using the eponymous Turing test, still well known and used today (Turing 1950). Turing’s work in turn influenced a new generation of AI in science fiction, from the sinister and disembodied *HAL 9000* in Arthur C. Clarke’s Space Odyssey series (1968), to the humanoid robot which passes the Turing test in *Ex Machina* (2014).

The popular fascination with AI and comparison to human intelligence has persisted throughout its history, often with dystopian predictions of AI ‘taking over’, fuelled by works such as Ray Kurzweil’s *The Singularity is Near,* which predicts that by 2029 Artificial General Intelligence (AGI) will match the intelligence of human beings (Kurzweil 2005). Others are much more sceptical and believe it will be centuries until AGI is reached, and the reality of AI has almost always disappointed in comparison to its representation in fiction. Whilst early AI aimed to replicate the human mind with AGI, today the focus has moved to Narrow AI, which is designed to complete specific tasks, and is more achievable with current technology. In practice, AI is developed in the realm of computer scientists, with only tenuous links to the cognitive sciences which inspired it. The aim of AI is now less in replicating human intelligence, than in developing a new computational or algorithmic intelligence, and there is a growing sense that AI may surpass our own, yet may not closely resemble it. As Max Tegmark suggests, what sets human intelligence apart is consciousness and emotion, rather than intelligence. He writes,

> *“As we humans prepare to be humbled by ever smarter machines, we take comfort mainly in being Homo Sentiens, not Homo Sapiens.”*
> Tegmark 2017

Tattersall, I. 2020. “Homo sapiens.” Encyclopedia Britannica. | Russell, S., Peter N. 2010. Artificial Intelligence: A Modern Approach. 3rd ed. Prentice Hall. | Akst, D. 2020, AI in Popular Culture: How Much Do You Remember?. https://www.wsj.com/articles/ai-in-popular-culture-how-much-do-you-remember-11604437200. | Turing, A. M. 1950. Computing machinery and intelligence. Mind, 59, 433–460. | Kurzweil, R. 2005. The Singularity is Near: When Humans Transcend Biology. New York: the Penguin Group. | Tegmark, Max. 2017. Life 3.0: being human in the age of artificial intelligence.

Gary Kasparov loses to Deep Blue: a landmark moment in the history of AI, 1997. Peter Morgan/Reuters.

The importance here is that whilst our intelligence is likely to be surpassed by machines, our consciousness will continue to set us apart. Perhaps surprisingly, given the anthropocentric nature of his Turing Test, Turing himself brought up this possibility as an objection to his own work, stating the question,

> *“May not machines carry out something which ought to be described as thinking but which is very different from what a man does?”*
> Turing 1950

Turing acknowledges this as a ‘very strong’ objection to his work, yet simply states that if a machine passes the Turing Test and is viewed as human, the question will become obsolete (Turing 1950). 70 years later, the Turing test has still not been passed, despite a number of attempts and unsubstantiated claims of success. At least for the time being, this is proving Turing’s objection correct - whilst there are intelligent algorithms today, they do not excel at the broad range of tasks that the human brain does.

It is in specific tasks that AI can get the upper hand and the classic way of testing this has been through games. As Alan Turing was exploring machine intelligence in the 1940s, he saw chess as a way to test his work, creating Turochamp with his colleague David Champernowne. This was an algorithm that could play a full game of chess and was too complex to run on computers of the time (Stezano 2017). Since then, chess became the highest challenge for AI researchers, and it wasn’t until Gary Kasparov’s loss to IBM supercomputer, Deep Blue, in 1997 that a world-champion chess player was beaten by a computer. Symbolically significant in the development of AI, this achievement has since been played down, and chess viewed as a game which can too easily be beaten with brute force calculations. The AI used relied on traditional rules and heuristics, written by humans and programmed into the machine to account for every possible move within the game. Today, a smartphone

Turing, A. M. 1950. Computing machinery and intelligence. Mind, 59, 433–460. | Stezano, M. 2017. In 1950, Alan Turing Created a Chess Computer Program That Prefigured A.I. https://www.history.com/news/in-1950-alan-turing-created-a-chess-computer-program-that-prefigured-a-i | Kasparov, G. 2018. ‘Chess, a Drosophila of Reasoning’. Science 362, no. 6419.

chess app can be stronger than Deep Blue (Kasparov 2018).

Following Deep Blue's victory, the new challenge for AI became the ancient game of Go. With more possible combinations of moves than there are atoms in the universe, Go represents a much more significant challenge and could not be tackled with the same brute force methods (Silver & Hassabis 2016). It was not until 2016 that the greatest Go player of decade, Lee Sedol, was beaten by AlphaGo. Developed by Google DeepMind, AlphaGo utilised two deep neural networks which played thousands of games against themselves in a trial and error process known as reinforcement learning, and was able to develop its own strategies based only on the rules of the game. The key here is that it required no formal training in past games or specific programming suited to Go, instead only using general machine learning technologies that can be applied to any rule based problem. In fact, a later iteration, AlphaZero, was able to teach itself to master Chess, Go and a number of other games. It became the strongest player in history, for each game, after only hours or days of reinforcement training, playing millions of games against itself in that time. Not only this, it came up with new strategies and moves which humans had never had the foresight to consider before, forming its own playing style and opening new possibilities for the games (Silver et al. 2018). Both Lee Sedol and Fan Hui, another Go player who competed in many games against AlphaGo over a period of months have claimed that their perception of the game has changed completely, and Hui's world ranking has risen greatly (Metz 2016).

What is clear from this example is that exposure to the new strategies of AlphaZero not only marked a development in AI, but also in intelligence generally. Human players are now utilising strategies conceived by an AI to reinvent the game of Go and improve their own performance. This signals a shift away from the pitting of human versus machine that began with the Turing Test, and reveals the potential of human-machine collaboration. Whilst machines can excel to super-human levels in specific tasks, it is the human that must define the tasks and interpret the results.

Perception

"An artificial neural network looks out on the world, trying to make sense of what it is seeing, in the context of what it has seen before. But it can only see the world through the filter of what it already knows, just like us."
Memo Akten, 2019

The technology of AlphaZero, broadly known as deep learning, and powered by artificial neural networks, offers the potential for strong AI of the future. It is used widely today from Instagram and Spotify recommendations to self-driving cars and facial recognition. In architecture, it's potential is largely unfulfilled, with only a small number of pioneering, 'neural architects', and academics exploring the field. The AlphaZero example illustrates how deep learning can provide innovative solutions to age old problems, overpowering centuries of human strategy and wisdom in a only a matter of hours. Clearly, when applied to well-defined problems such as games, this technology can provide genuine intelligence far beyond what can be achieved by the human brain alone. But is this form of intelligence so different to the way our own minds work?

When discussing intelligence, it is important to consider the language used and why we are using it. Deep neural networks (DNNs) are loosely based on our understanding of human intelligence from the field of neuroscience, specifically the structure and hierarchy of the human visual cortex (Campo and Manninger 2019). In this field,

Silver, D, Hassabis, D. 2016 AlphaGo: Mastering the ancient game of Go with Machine Learning. Google DeepMind. https://ai.googleblog.com/2016/01/alphago-mastering-ancient-game-of-go.html. | Silver, D, et al. 2018. AlphaZero: Shedding new light on chess, shogi, and Go. https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go. | Metz, C. 2016. In Two Moves, AlphaGo and Lee Sedol Redefined the Future. Wired. https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/. | Memo Akten et al. 2019. Learning to see: you are what you see. SIGGRAPH 2019. Association for Computing Machinery, New York, Article 13, 1–6.

intelligence is defined as the ability to learn from experience and to adapt to, shape, and select environments (Sternberg 2012). Similarly, Russell and Norvig (2010) state that a truly intelligent artificial agent must have the ability to learn from and interpret experience to meet goals adaptively. The first step in intelligence is therefore the ability to learn, and this must happen through experience. As humans, we experience the world around us through our perception, and DNNs experience in much the same way, building up their own perception based on the data they are given.

Of particular interest to architects and designers are Generative Adversarial Networks (GANs), due to their visual nature and generative capabilities. GANs are composed of two competing deep convolutional neural networks (CNNs), a generator network and a discriminator network. The generator produces an image from a random field of noise, whilst the discriminator learns to distinguish whether the image is a member of the data set, or not. As the process repeats, the generator learns to synthesise images that cannot be distinguished from the original data (figure 1).

Through a process of analysis and pattern recognition of large data sets, GANs are able to build their own perception. By interpolating semantic characteristics of the training data, they learn to hallucinate outcomes which are entirely novel, yet familiar. This concept of machine perception and hallucination can be compared to our own human perception. When Anil Seth explains that, 'your brain hallucinates your conscious reality', he shows that our perception is not a direct view of the world, but is a prediction based on sensory input and our past experience (Seth 2017). The electrical signals, or data, our brain receives from our senses is open to interpretation, and we are only able to recognise and identify objects, or spaces, through our prior knowledge of them. As architects, our perception is shaped by our education and environment. We are trained in certain ways depending on the school we attended and the cultural, political and social context in which we live. This forms certain 'schools of thought', and each architect can be said to have their own sensibility. In much the same way, the perception

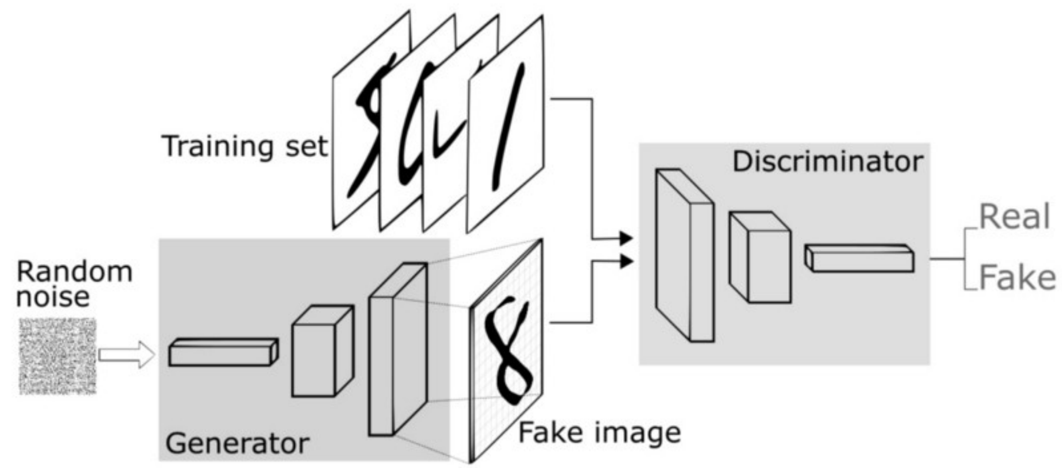

Figure 1: GAN Typical Architecture (Chaillou 2019).

Campo, M., and Manninger, S. 2019. "Imaginary Plans." Proceedings of the 2019 ACADIA Conference - Ubiquity and Autonomy. | Russell, S., Peter N. 2010. Artificial Intelligence: A Modern Approach. 3rd ed. Prentice Hall. | Chaillou, S. 2019. 'The Advent of Architectural AI'. Harvard Graduate School of Design. | Sternberg R. J. 2012. Intelligence. Dialogues in clinical neuroscience, 14(1), 19–27. | Seth, Anil. 2017. "From Unconscious Inference to the Beholder's Share: Predictive Perception and Human Experience." PsyArXiv.

of a neural network depends entirely on the training data given to it.

This comparison between human and machine perception can help to reflect on our own biases and understand those which can be developed by neural networks. When we design, we can never be completely objective, and there is always an underlying aim or preference. This raises a number of questions over the potential of neural networks in design. Can deep learning allow us to produce the unexpected? Can it allow us to remove our own filters and preconceptions and allow us to perceive without bias? Or conversely, can a network learn our sensibility and design to satisfy our preferences?

Creativity

The questions asked above ultimately lead back to questions of creativity. Can a machine be creative? This is the most common question that arises with the use of generative neural networks, but as Zylinska (2020) suggests, perhaps this is not the best question to be asking. She instead points towards philosopher of technology Vilém Flusser, who, when considering photography, suggested that human and machine agency are in constant entanglement. Flusser states that,

”This is a new kind of function in which human beings are neither the constant nor the variable but in which human beings and apparatus merge into a unity”.
Flusser 2000

This suggests the creativity of a machine cannot be separated from that of the human, and in explaining either we must first define creativity itself. Whilst there have been many disputed definitions of creativity, Kuszewski (2009) sets out the key to creativity as, “the concept of generating novel ideas that are appropriate to the situation at hand”. To be creative a person must not only generate novel ideas, but these ideas must be relevant and serve a useful purpose in that situation. Du Sautoy (2019) takes a similar position, summarising the theory of creativity as, “the drive to come up with something that is new, that is surprising, and that has value.”

For machines, coming up with something new is easy, neural networks can synthesise thousands of novel images very quickly. But how are surprise and value measured? Mitchell (2019) states that, “being creative entails being able to understand and judge what one has created”. The key to recognising both surprise and value is judgement, it is not possible to be surprised without first having an expectation of what will occur. Creativity must therefore come from producing the unexpected, from breaking free of our preconceptions and patterns of thought. But how does a person do this? How does a human create something that it cannot expect? Post-humanist art theory would suggest that, throughout history, all art works produced by humans have been the result of interaction and experimentation with an abundance of non-human agents. From organic objects, apparatus and drugs, to contemporary technologies, computers and AI (Zylinska 2020). The human is inherently moulded by its interactions with non-humans, as explained by Culkin (1967),

”We become what we behold. We shape our tools and then our tools shape us”.
Culkin 1967

If human technology began with the creation of the first stone tools, then this creative act began Flusser’s cyborg vision of a merger between human and apparatus. Experimentation with non-human agents can allow us to produce novel ideas that break free of our expectations. Contemporary technologies, such as neural networks, can create an abundance of novelty, but this alone does not constitute creativity. Only the human can provide the judgement required to determine valuable and surprising results. To maximise creativity, we must embrace our cyborg nature. The focus must be on a unified process between human and technology, as Bolojan (2021) states in his exploration of neural networks,

“The interest should be in the creativity of the feedback loop between human and machine - not just in the machine itself, and not just in the human.”
Bolojan 2021

Zylinska, J. 2020. AI Art: Machine Visions and Warped Dreams. London: Open Humanities Press. | Flusser, Vilém. 2000. Towards a Philosophy of Photography. London: Reaktion Books. | Kuszewski, A. 2009. The Genetics of Creativity: A Serendipitous Assemblage of Madness. Metodo. | Du Sautoy, M. 2019. The Creativity Code: Art and Innovation in the Age of AI. UK: Harvard University Press. | Mitchell, M. 2019. Artificial Intelligence: A Guide for Thinking Humans. Penguin UK. | Culkin, J. 1967. A Schoolman’s Guide to Marshall McLuhan. Saturday Review.

Intention and Methodology

Based on the proceeding analysis of machine intelligence, perception, and creativity, this thesis proposes to investigate the following questions:

1. **Can neural networks be used to challenge our own preconceptions?**

2. **How can neural networks be used to augment the creative process of architectural design?**

3. **How can architects best prepare datasets for use with deep learning?**

These questions will be addressed using two case studies of architects who are exploring similar topics in their work. These case studies will then inform my own series of experiments which will explore the curation and design of custom datasets and their use with deep learning neural networks.

Bolojan, D. 2021. FIU DDES Lectures_Theories of the Digital_Session4: AI and Creativity. Youtube.

Case Study 1: Serpentine2020?

Serpentine2020? is a project developed by Immanuel Koh that explores a neural networks understanding of a designers sensibility. The project was initiated after the 2020 Serpentine pavilion was scrapped due to the Covid-19 pandemic. It questions whether a neural network can synthesise a new pavilion for 2020, based on 20 pavilions from previous years.

Koh explains that as an architectural student you study and learn from these important and widely publicised buildings, forming your own interpretation and perception which you use to develop your own work (Koh 2021). Can neural networks do the same thing? He goes on to question whether the characteristics of the individual architects and pavilions remain in the synthesised images, in the same way that precedents and influences can be evident in an architect's work.

The synthesised images were created using a GAN and display a range textures, forms and colours that are novel yet reminiscent of the original pavilions. Our human perception can recognise these influences in the synthesised images as we associate particular characteristics to certain architects. For example when we see the red forms of figure 3 we immediately recognise Jean Nouvel's pavilion from 2010 as the primary influence. In other images the GAN appears to have created hybrids that mix features from different architects together.

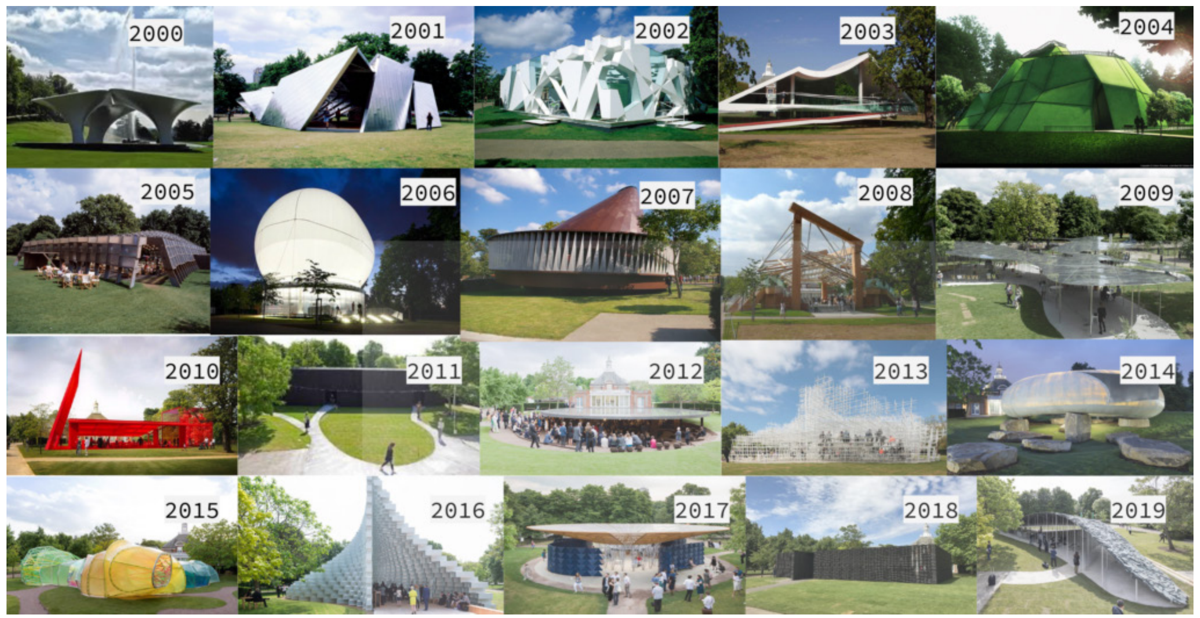

Figure 1: Serpentine Pavilion 2020?: Training set using images of 20 previous pavilions. Immanuel Koh, 2020.

Koh, I. (2020). Serpentine2020?. https://artificial-architecture.ai/?p=140. | Koh, I. (2021). Serpentine2020?. FIU DDES Lectures: Theories of the Digital Session 6 on AI & Architectural Design. https://www.youtube.com/watch?v=u-8KO_7ycgE

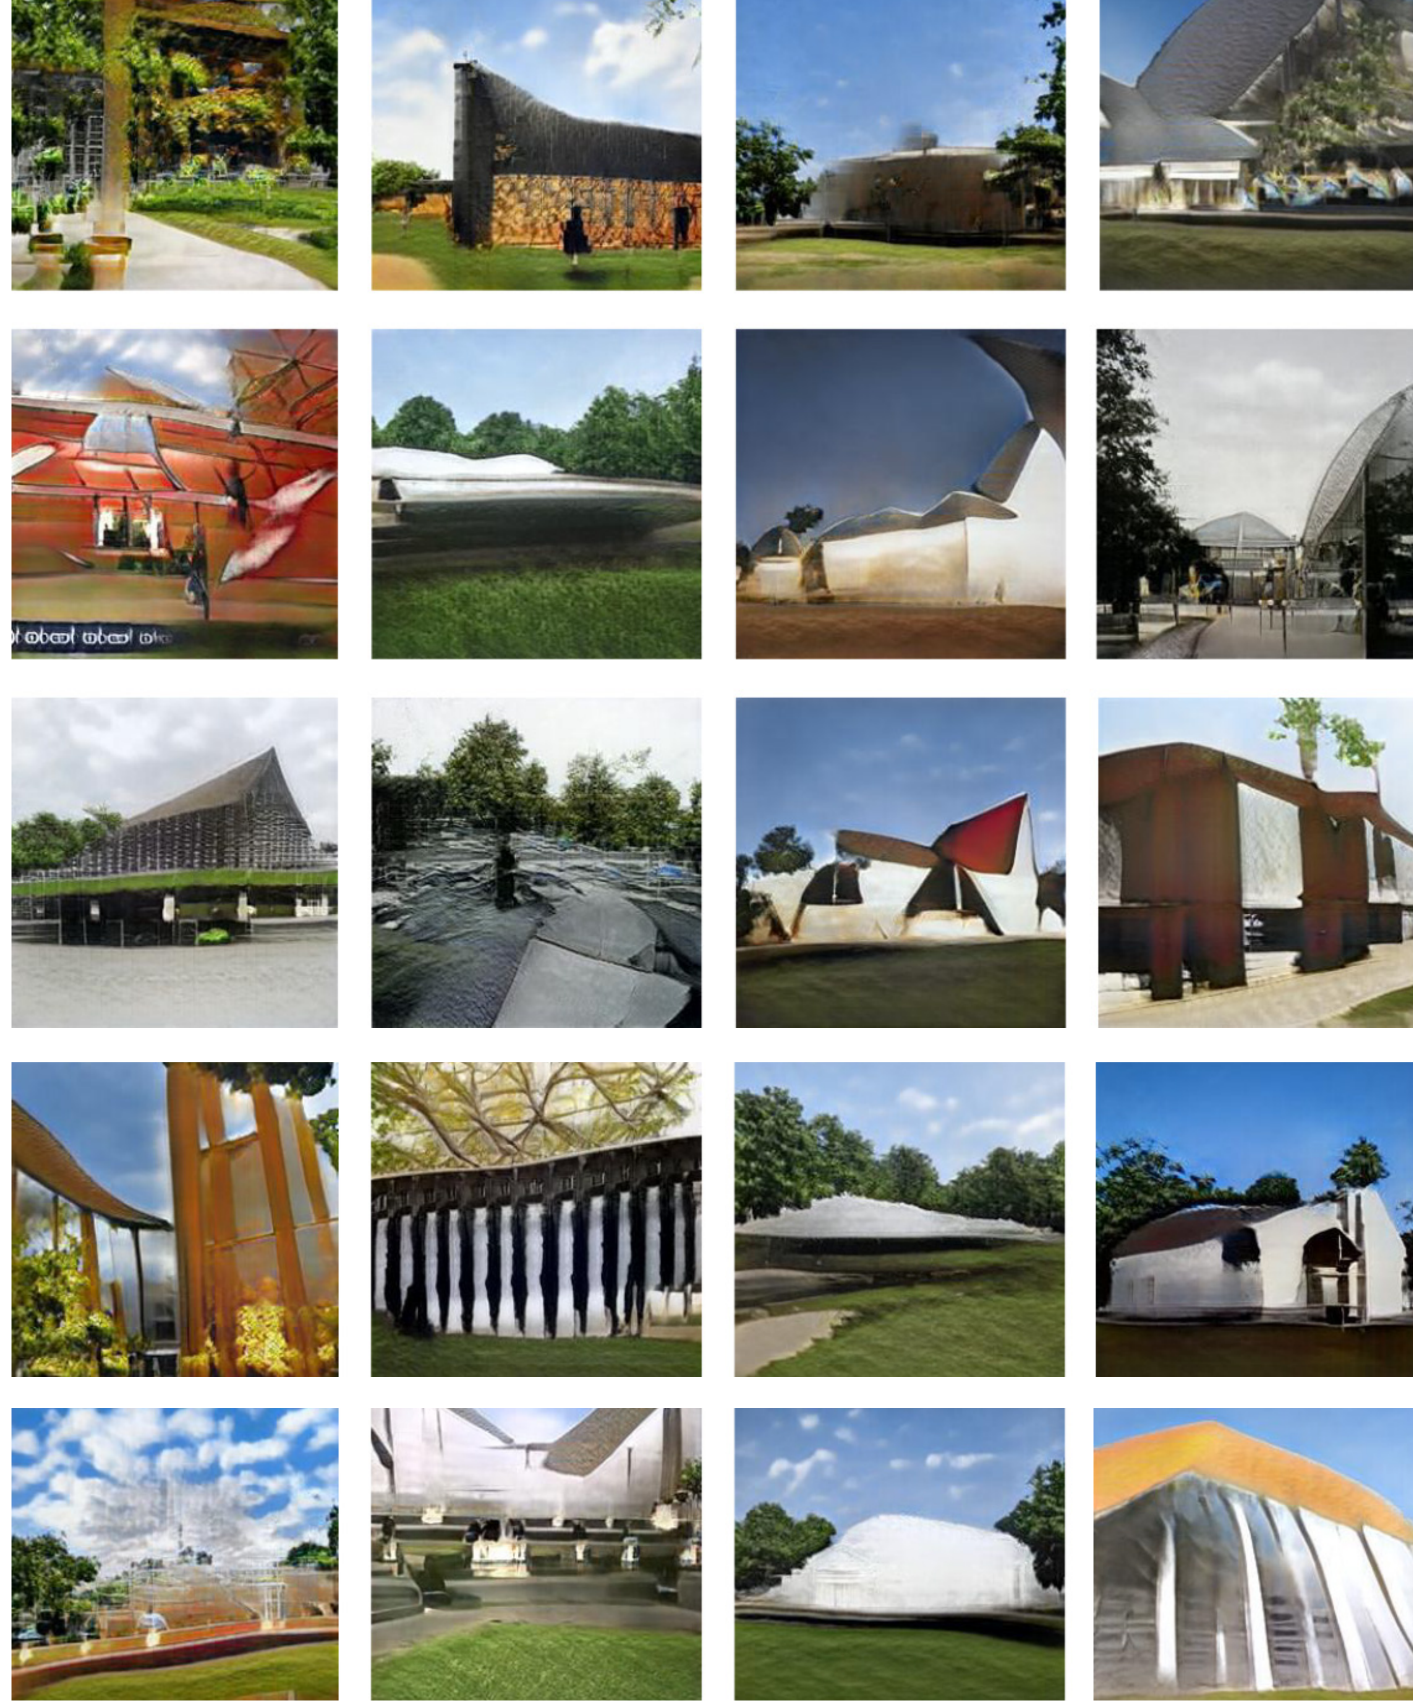
Figure 2: Serpentine Pavilion 2020?: Synthesised Images. Immanuel Koh, 2020.

The second part of Koh's experiment tests whether a machine can classify these same influences. The synthesised images are plugged into a second CNN, known as an EfficientNet Classification model, which can predict the percentage of each pavilions influence present in each image. These results are shown in figure 4, with the opacity of the image representing the strength of influence. Whilst there is no specific criteria on which to judge this classification, the results appear to follow my personal perception of them to varying degrees. Output Image 33 is clearly recognisable as an amalgamation of the input images shown. Least recognisable is output image 76 which the model has stated is 100% influenced by the Frank Gehry pavilion of 2008. This is perhaps a formal relationship, although to my eye the pavilions of 2005 and 2007 are a closer match in colour and form.

Figure 3: Serpentine2020?: Synthesised Image reminiscent of Jean Nouvel. Immanuel Koh, 2021.

Analysis

This project has an interesting premise of testing the perception of neural networks to interpolate and distinguish between the work of specific architects, however the results are quite limited and difficult to evaluate. One building from each architect is not enough to determine the architects preferences and sensibility. In order to understand the consistent characteristics present the neural network would need many examples from each architect. Even so, the model is able to successfully synthesise novel images with characteristics such as form and colour inherited from the training data, and classify the synthesised images with some success. The experiment has shown the potential of GANs to iterate design possibilities which could be used for creative inspiration in the design process. The low resolution of the images leaves room for open-ended interpretations which could lead to varied design proposals. It has also shown significant parallels between human and machine perception by creating classifications that match my intuitive human judgement.

The project uses a dataset of images of past pavilions that has been assembled from images on the web. This process is known as data scraping and raises questions over data privacy and copyright. When building my own datasets this is something to be wary of, and will be explored in my later design experiments.

Koh, I. (2021). Serpentine2020?. FIU DDES Lectures: Theories of the Digital Session 6 on AI & Architectural Design. https://www.youtube.com/watch?v=u-8KO_7ycgE

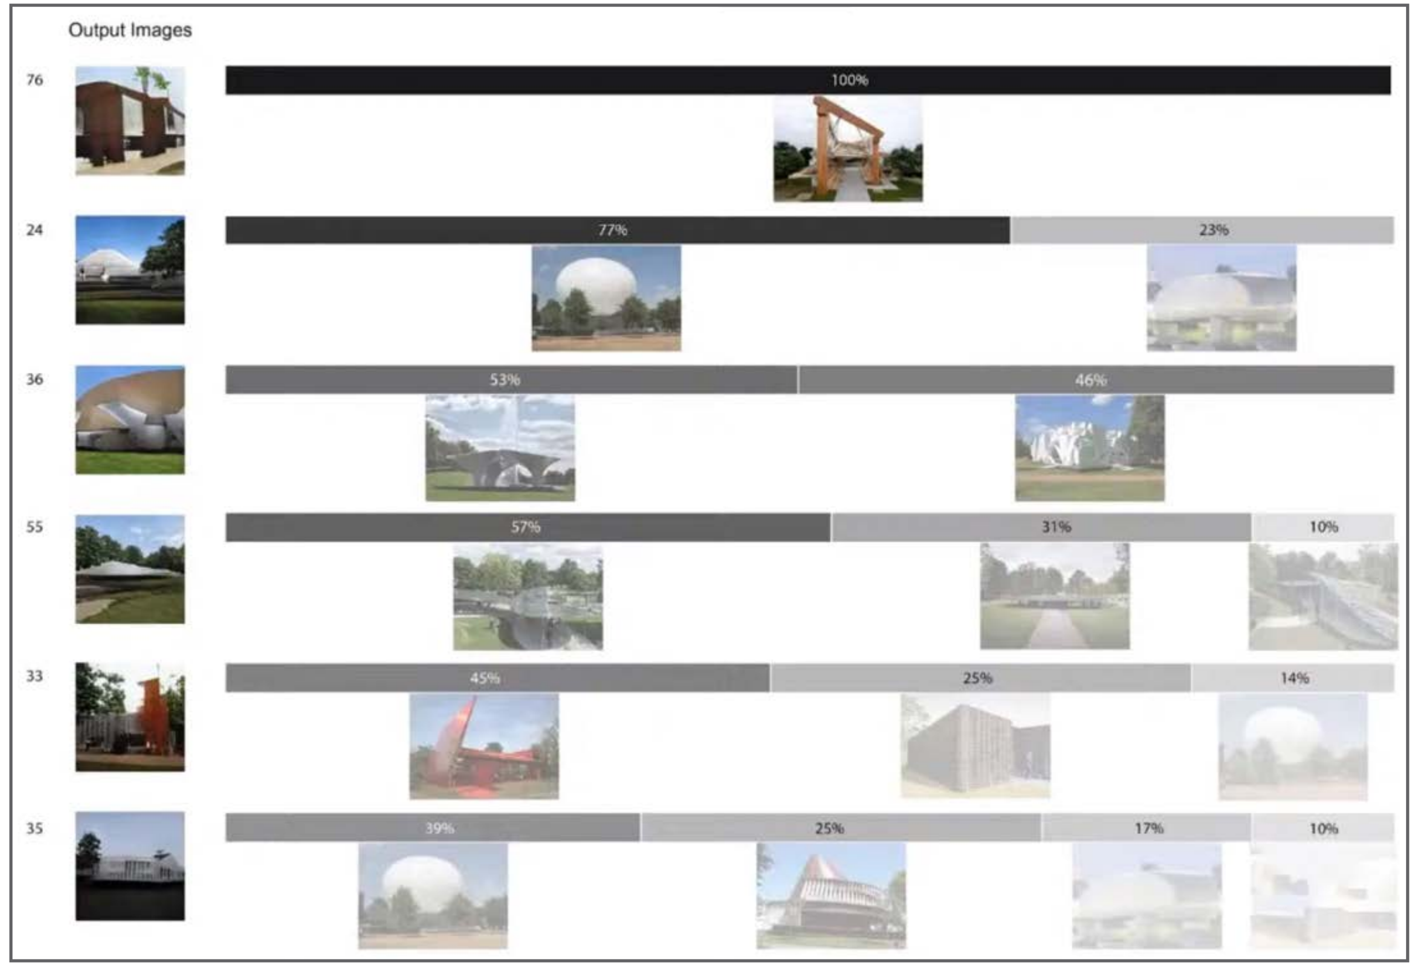

Figure 4: Serpentine2020?: Percentage influence of original pavilions in synthesised images, Immanuel Koh, 2021.

Case Study 2: DeepHimmelb(l)au

"DeepHimmelb(l)au is designed to interact with designers and inspire creativity. It is designed to facilitate a medium of constant interaction / feedback loops between designer interpretations and its own interpretations, between designer perceptions and its own perceptions. In that sense we are aiming – through augmentation – to strengthen our capabilities as creators – a collaboration between machines and humans."
Daniel Bolojan, 2020

This second case study continues the investigation into machine perception and creativity, but rather than comparing the designs of multiple architects, it explores the potential of GANs to learn and replicate a particular architecture firm's semantic characteristics. This has the potential application to be used as a design tool within the firm to augment the design work-flow and the architect's creativity.

Deep Himmelb(l)au, developed by Daniel Bolojan at Coop Himmelb(l)au (CHBL), is a project that operates at the intersection between architectural research, practice and AI. CHBL takes the stance that AI is a technological tool that should be used to augment the creative process of the designer. It tests the similarities and differences between human and machine perception and questions whether machines can discover perceptual deficiencies in human recognition. Do we as humans have perceptual blind spots that machines can uncover? This brings us back to the questions of bias and preconceptions in perception. As humans we have certain expectations that effect our perception, and as architects specifically we are taught in certain ways. When we see particular drawings or images we instantly recognise them as plans or sections or perspectives, but a neural network only sees colour and composition. Can this help to see things in new ways and challenge our perspective? Our human perception is able to consciously, or unconsciously, reinterpret semantic representations from other domains into the domain of architecture. Can neural networks work in a similar way? (Bolojan 2020).

To address these questions CHBL approaches the use of AI with a number of methods, analysing the studio's design process to find areas in which AI can provide augmentation. This includes optimisation methods, though at the current stage is mainly focussed on the creative process. Deep Himmelb(l)au itself is a complex neural network that borrows from a number of well known networks, such as Pix2Pix, CycleGAN and Generative Query networks (Bolojan 2021). It is trained on CHBLs full repertoire of works which have been mapped by similarity of semantic characteristics (figure 1).

The network complexity and results have been built up and improved over a number of years and it has been tested in many different circumstances and stages of the design process. It was initially inspired by Pix2Pix works which translated label maps and sketches into images and its early use explored these topics. Figure 2 shows label maps being translated into rendered images. The next experiment then used the network to translate photographs and videos of physical foam models into rendered images, creating an interesting relationship between the physical and digital design space (Figure 3).

Bolojan, D. 2020. Meet Deep Himmelb(l)au. https://nonstandardstudio.com/2020/01/26/meet-deephimmelblau/. | Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

Figure 1: Deep Himmelb(l)au: Complete mapping of CHBL designs by similarity. Bolojan, 2021.

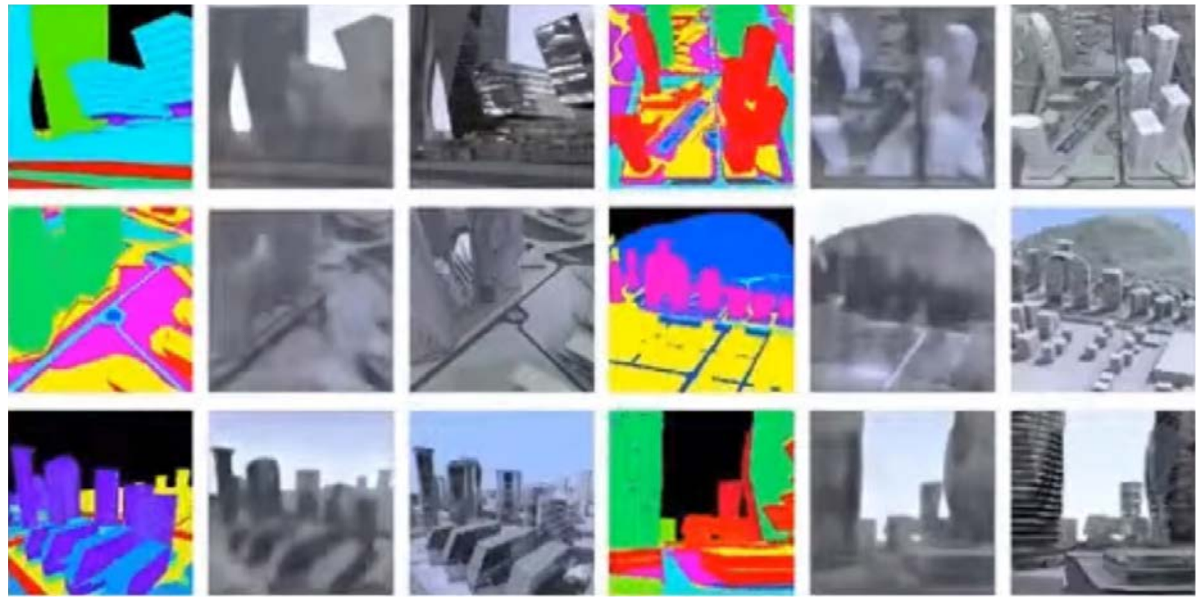

Figure 2: Deep Himmelb(l)au: Translating label maps to rendered images. Bolojan, 2021.

Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

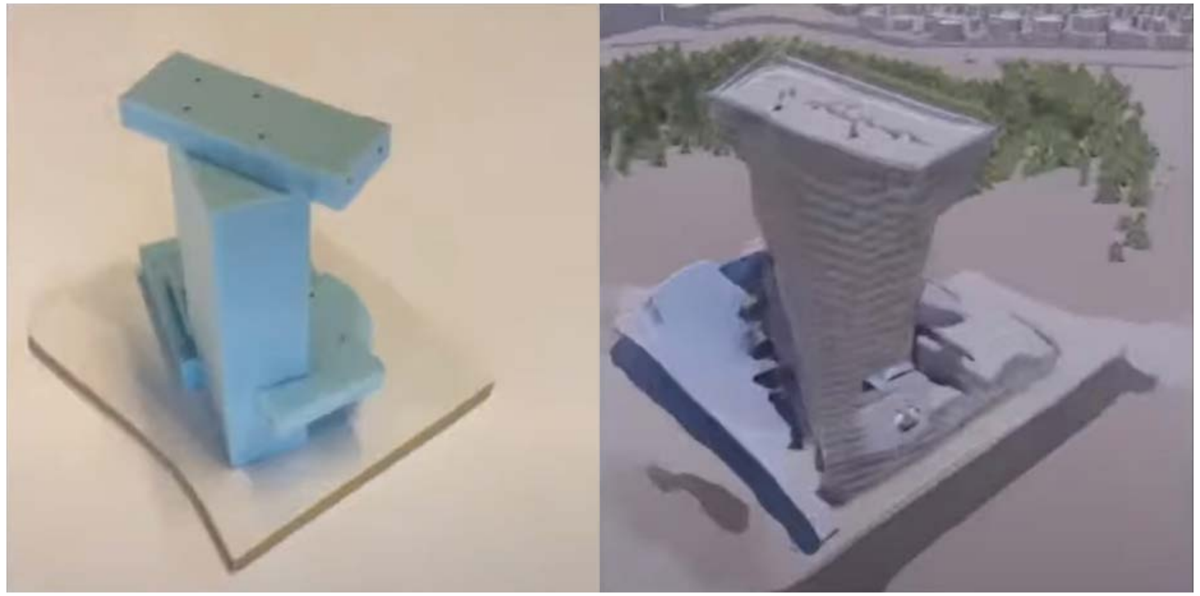

Figure 3: Deep Himmelb(l)au: Translating physical models to rendered images. Bolojan, 2021.

These early examples are still of low resolution and show a direct translation from physical to digital model.

Figure 4 shows a significantly improved resolution, and critically challenges our human perception of the physical model. We can clearly see a bird's eye view of a building with long ramps leading up to it, yet the neural network exhibits a completely different perspective. It appears to view the model as a facade, and by looking only at the composition of the input image, is able to generate a very unexpected and interesting result. The network is starting to open up new avenues for inspiration and creativity in the design process, creating a feedback loop where a human model could be interpreted by the network then reinterpreted by the human into a new design.

Figure 5 displays the next development in the network, where it now has some understanding of the image perspective, yet also adds facade detail and structures. Then in a reversal of the process the network is asked to translate from an image into a foam model. The image supplied is an image of a built CHBL building, yet once again the network interprets the image completely differently to how we perceive it. It has generated a semi-urban cluster of buildings with the composition of the original facade. There is a clear novelty and creativity in this process of human and machine feedback (figure 6).

The latest stages of Deep Himmelb(l)au are focussed on image generation without a specified input image, instead exploring the complete latent space of the network model. In this mode of image generation the images are generated from a gaussian noise input, which is mapped to a latent space model, such as a 100-dimensional hypersphere (Brownlee 2020). This allows the designer to cycle through the nearest points in the latent space to generate a series of synthesised images with a smooth transition between them. Figure 7 shows 100,000 synthesised

Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

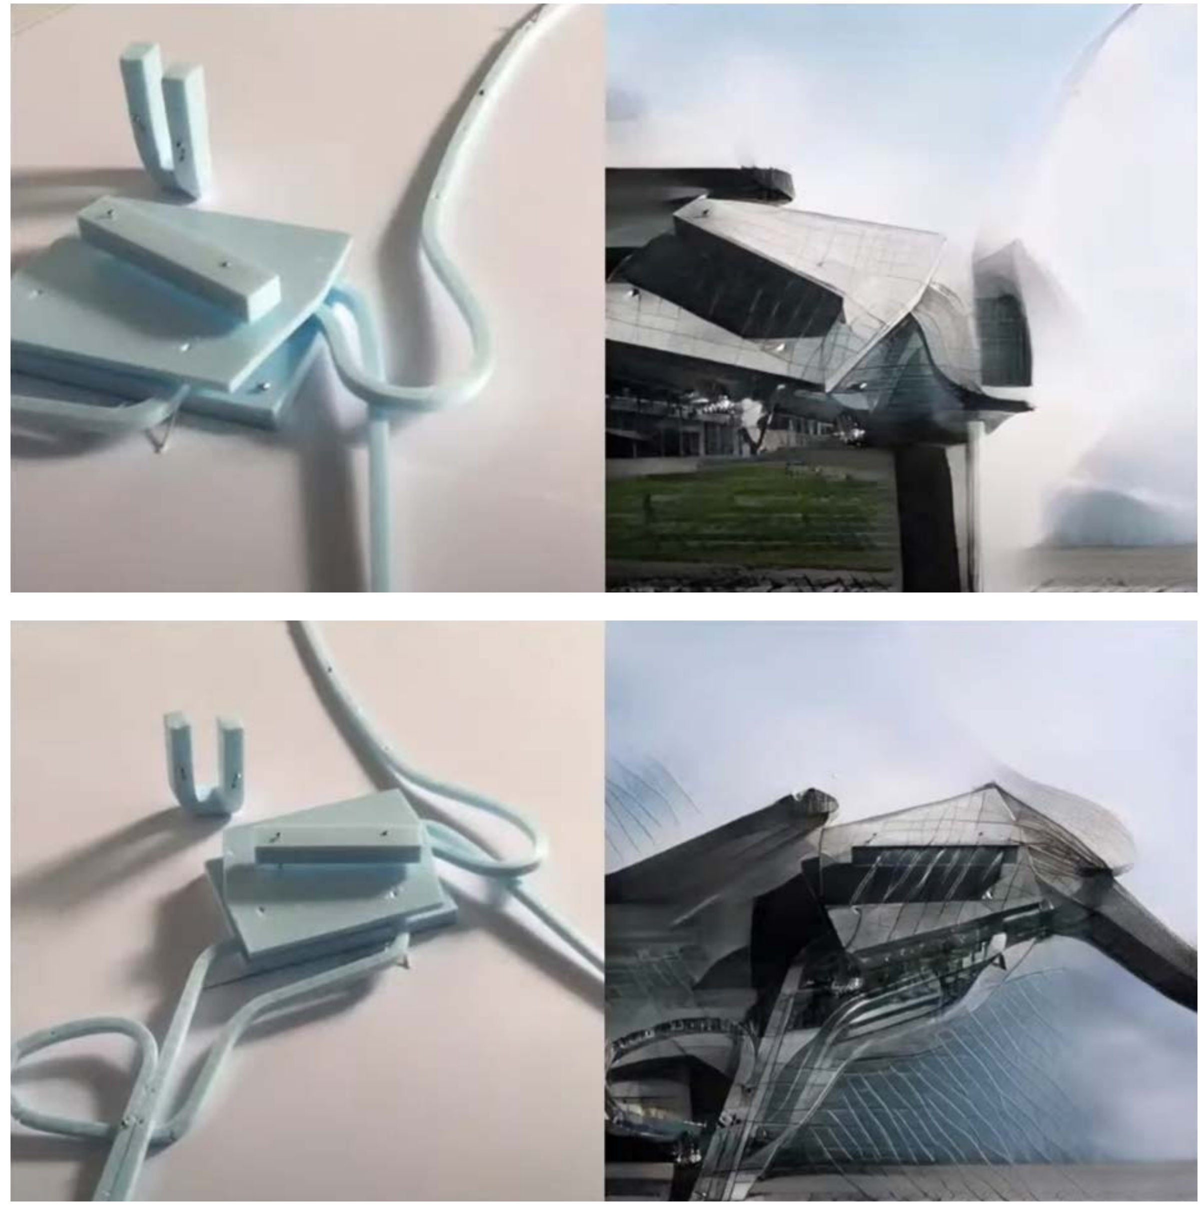

Figure 4: Deep Himmelb(l)au: Translating physical models into rendered images. Bolojan, 2021.

Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

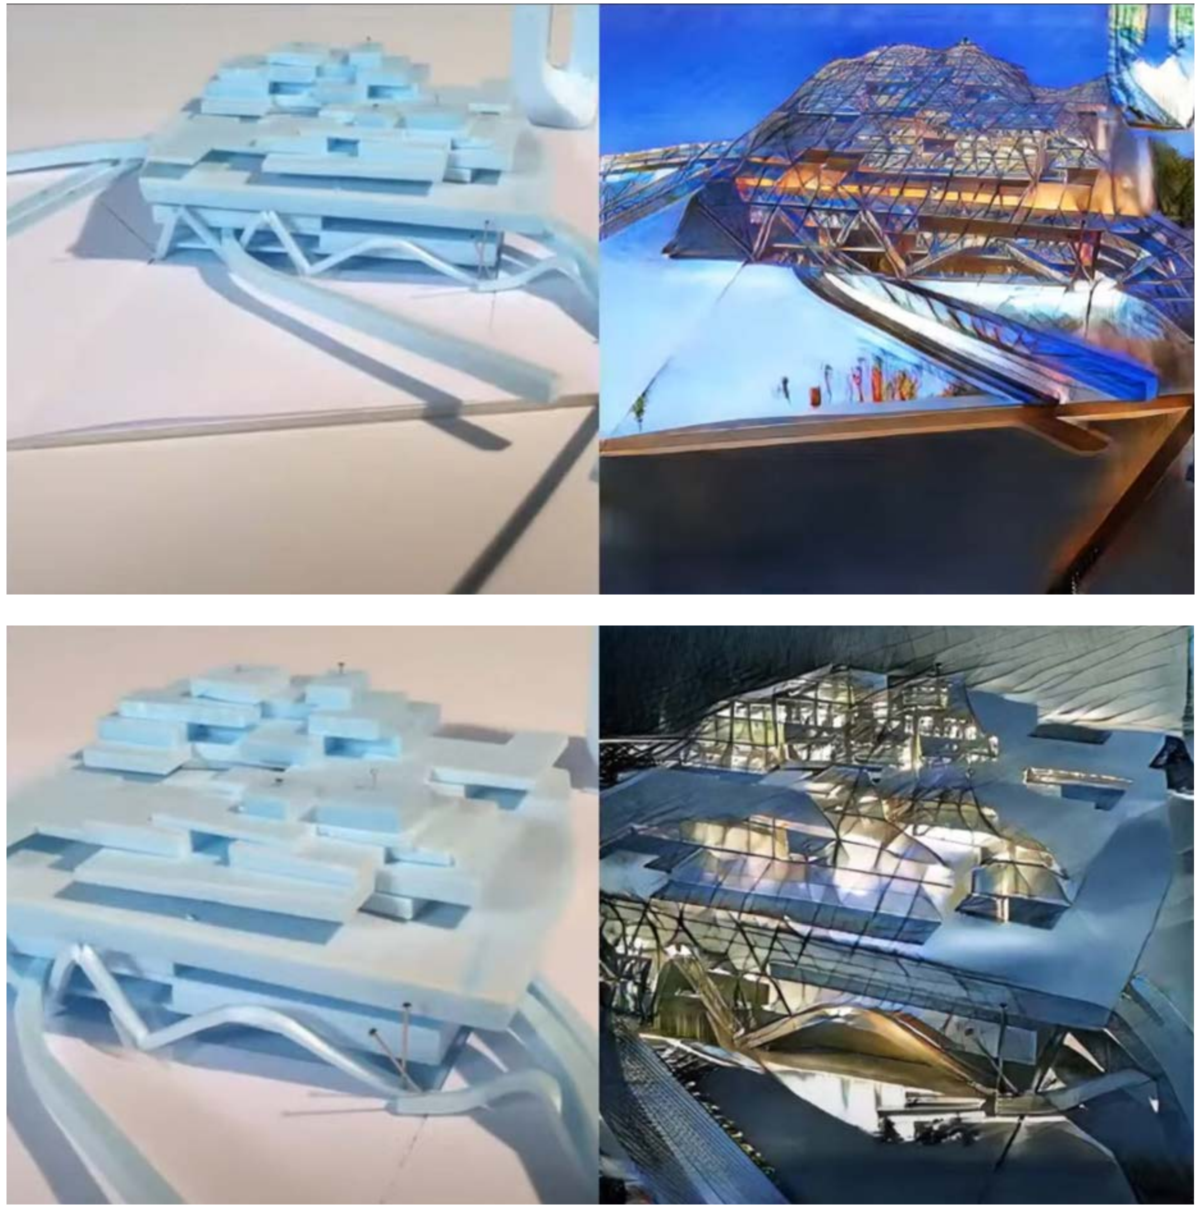

Figure 5: Deep Himmelb(l)au: Translating a foam model into a rendered image. Bolojan, 2021.

Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

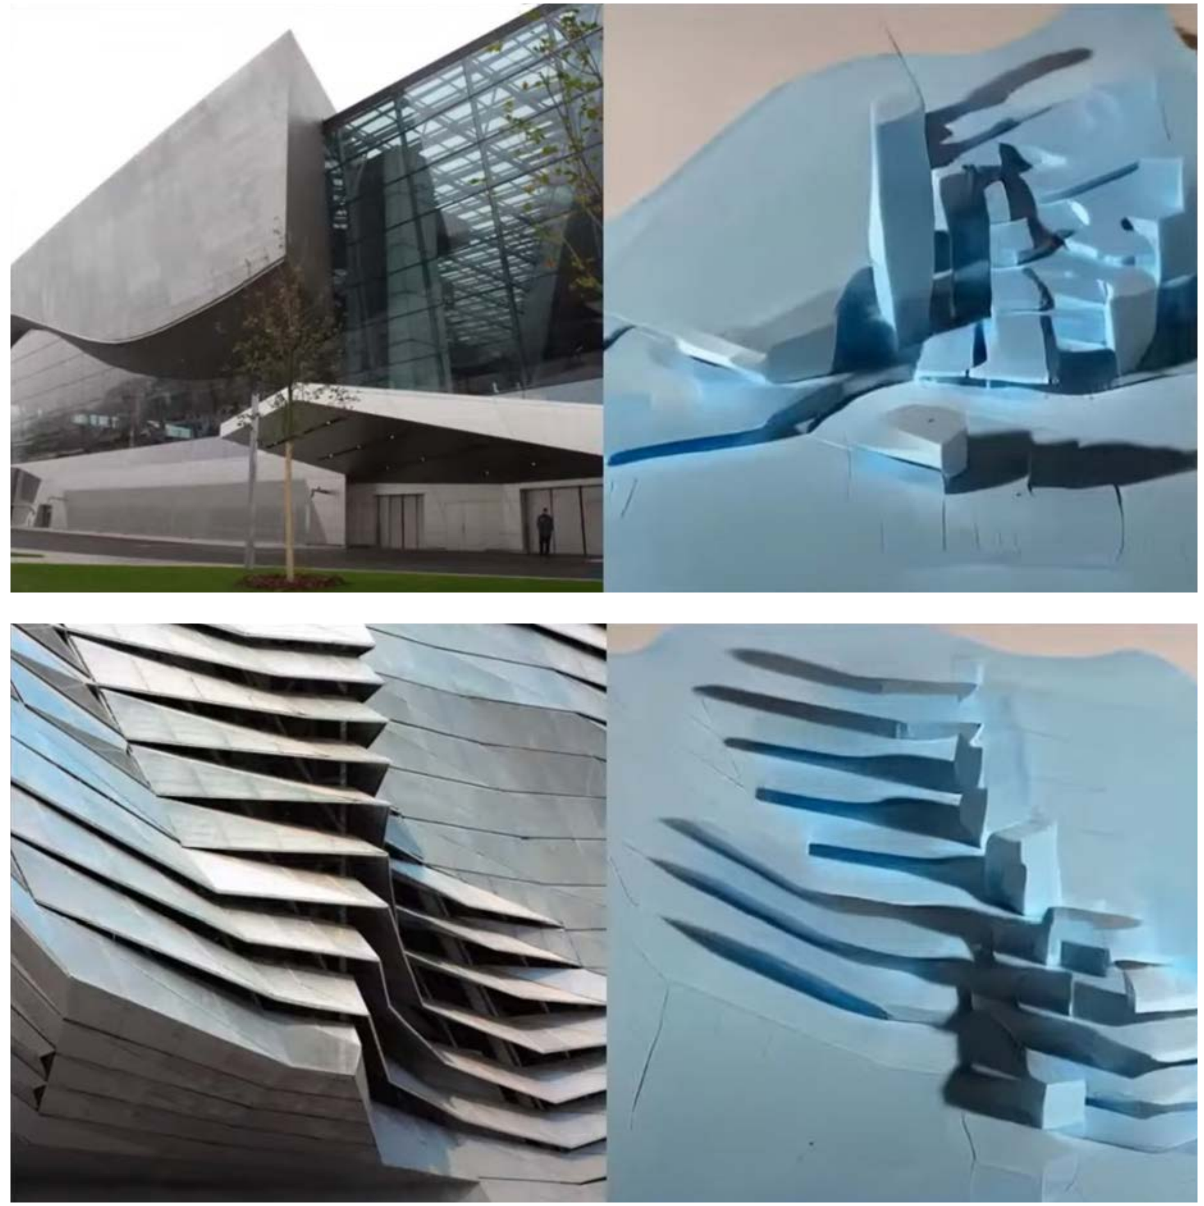

Figure 6: Deep Himmelb(l)au: Translating a built facade into a foam model. Bolojan, 2021.

Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

images mapped to the latent space. This allows for a single seed to be selected and searched to find similar images, with more or less variation determined by distance between the points in the latent space. This allows the designer to focus on a particular series of design options, exploring the model and refining their choice (figure 9).

Additionally, to improve the output of the model, additional layers of information were added to the training set, such as alpha channels, depth maps and material IDs (figure 8). These extra details, as well as adjustments to the network itself, and a higher resolution of training images, led to near photorealistic synthesised images (figures 11 & 12).

Analysis

Coop Himmelb(l)au, led by Wolf Prix, is one of the founders of the deconstructivist movement in architecture, that arose under the influence of philosopher Jacques Derrida. The deconstructivist movement is founded on the principles of unconscious processes, which challenge constructed rules and preconceived ways of thinking (Hoteit 2015). In architecture, the movement has led to widespread experimentation and a disturbance to the pure forms of modernism. Using AI, the practice has found a new, highly effective, method of designing with unconscious processes.

DeepHimmelb(l)au is the most advanced research undertaken into the design potential of neural networks in the architectural field. By exploring machine perception it is able to challenge current design processes and modes of creativity. In common practice, a designer is constantly and consciously judging their work to preconceived standards and preferences. The use of neural networks can allow the designer to break free of their own expectations, challenging their own perception and widening the possibilities of creation.

The images synthesised by the highly developed

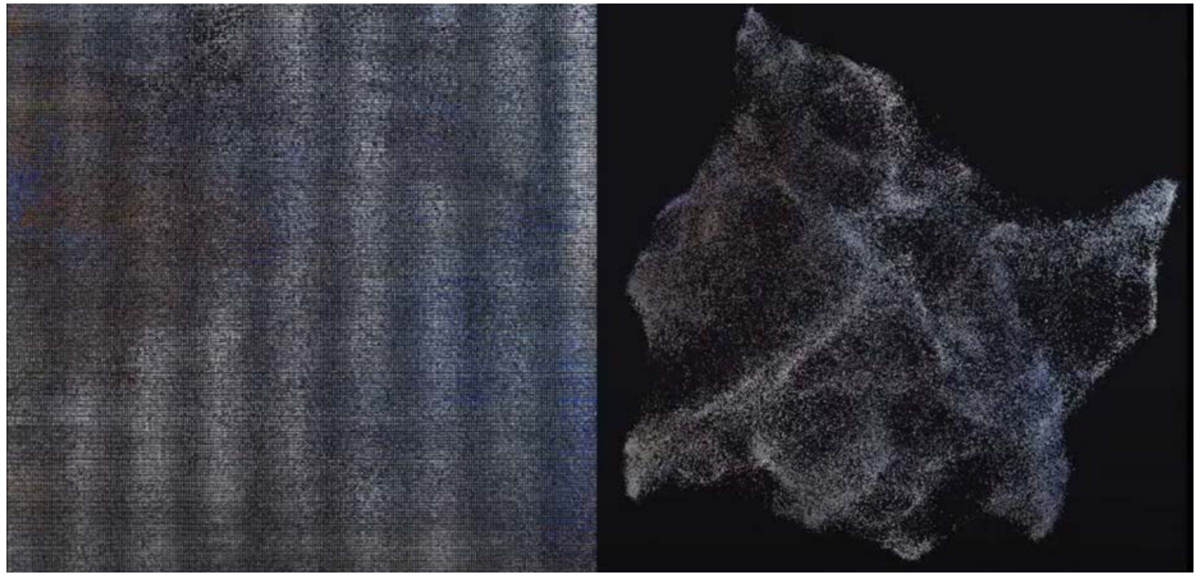

Figure 7: Deep Himmelb(l)au: 100,000 points mapped in the latent space. Bolojan, 2021.

Brownlee, J. 2020. How to Explore the GAN Latent Space When Generating Faces. Machine Learning Mastery. | Hoteit, A. 2015. Deconstructivism: Translation From Philosophy to Architecture. Canadian Social Science, 11(7), 117-129. | Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

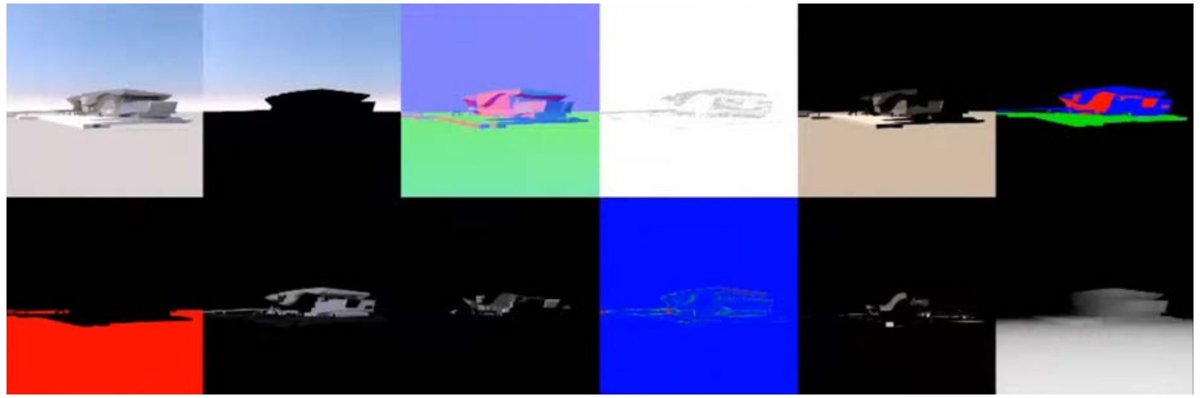

Figure 8: Deep Himmelb(l)au: Additional layers of information used to train the network. Bolojan, 2021.

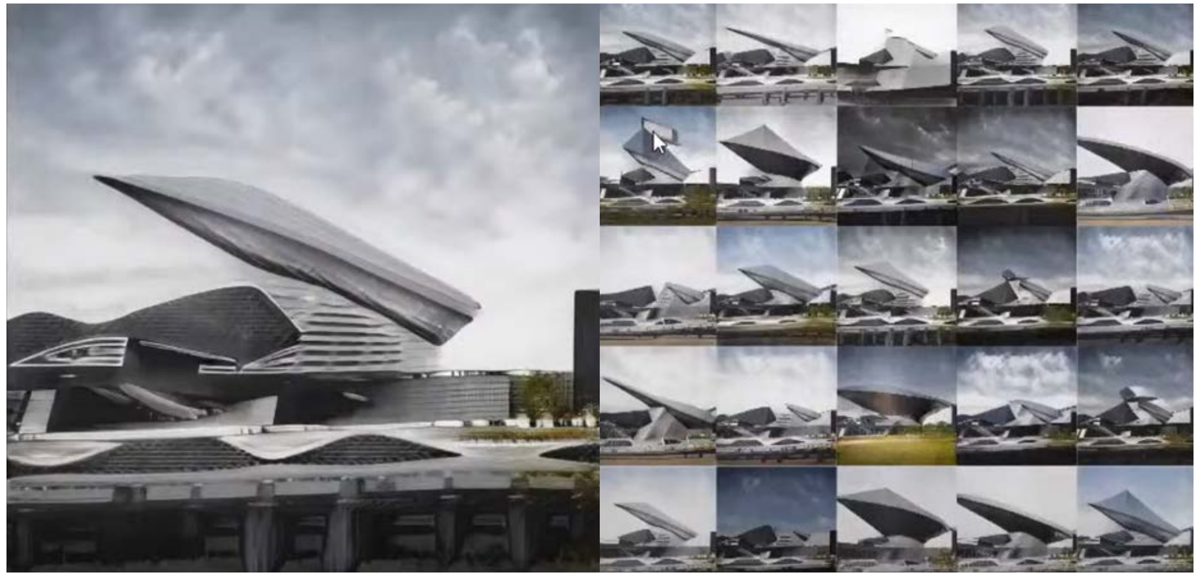

Figure 9: Deep Himmelb(l)au: A seed image is searched to find nearby points in the latent space. Bolojan, 2021.

Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

DeepHimmelb(l)au model are almost photorealistic in quality, and whilst novel in form and composition are quite clearly recognisable as CHBL buildings. The network has developed a perception that captures the design sensibility of the studio through interpolation of its past work. This is effective due to the large body of work the firm has built up over many years of practice, and its consistent characteristics. However, this consistency and use of past work as training data raises its own questions. By training the GAN using only your own work, are you not simply encoding your own biases into the neural network?

O'Neil (2017) states that, "algorithms are opinions embedded in code", and that they only repeat our past practices. By choosing the data and the definition of success we are injecting our own bias into algorithms. Whilst in many fields this can lead to unbalanced and discriminatory algorithms, in CHBL's case, it seems that this is exactly the aim. By using the neural network to replicate their own design sensibility, CHBL is purposefully training it to be biased towards certain forms, patterns and finishes. When using datasets of architectural façades or renderings, this is not necessarily a problem, but may limit design possibilities. If architects begin to use other types of data, for example building occupancy or client data, then these issues of bias will become hugely important to consider and overcome.

When considering how to build my own datasets, the question arises of how this technology can be used by designers or firms, such as myself, who do not have a body of work on which to train a network, or, do not want to limit themselves to training data based on their previous work. This question will be explored in my design experiments, in which I investigate the curation and generation of custom datasets for use with deep learning.

Figure 10: Deep Himmelb(l)au: Synthesised images, early model. Bolojan, 2020.

O'Neil, C. 2017. The Era of Blind Faith in Big Data Must End. TED Talks. https://www.ted.com/talks/cathy_o_neil_the_era_of_blind_faith_in_big_data_must_end/. | Bolojan, D. 2020. Meet Deep Himmelb(l)au. https://nonstandardstudio.com/2020/01/26/meet-deephimmelblau/.

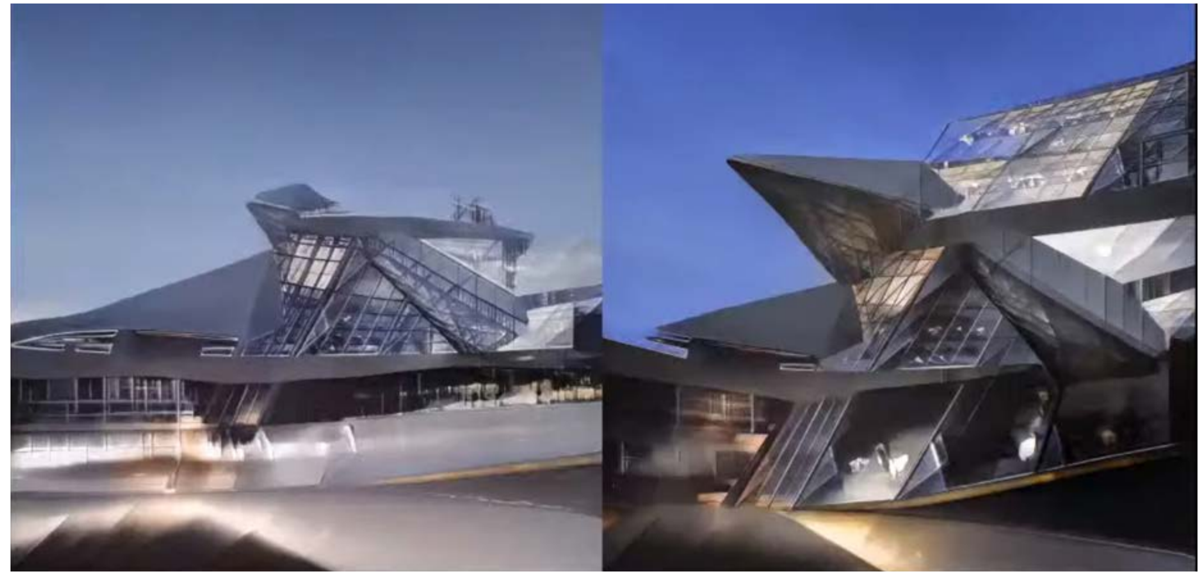

Figure 11: Deep Himmelb(l)au: Synthesised images, developed model. Bolojan, 2021.

Figure 12: Deep Himmelb(l)au: Synthesised images, developed model. Bolojan, 2021.

Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.

Design Experiments

The proceeding case studies have shown that neural networks are capable of challenging human perception in design, as well as capturing a designer's sensibility. These capabilities can stimulate creativity through a process of constant feedback and human-machine collaboration. Whilst Serpentine2020? was limited to a small dataset, it was able to show with some success that the neural network could perceive differences in architectural form and style. DeepHimmelb(l)au, trained on a much larger and more cohesive dataset, showed that a network could be trained to highly effectively produce novel forms and compositions in keeping with a firm's sensibility.

The two case studies have exhibited two types of dataset, the first scraped from the internet and the second using a studio's own library of past work. Each type has its own benefits, scraped datasets offer a wide range of possible images, whilst a library of past work can contain highly specific data, and can be augmented with additional image or data layers.

When considering how to create datasets for my own experiments, both methods also have drawbacks. Firstly, scraped datasets have issues of ownership and copyright which must taken into account, and secondly, many young designers and firms, such as myself, will not have a catalogue of previous built work or digital models that they can use. Additionally, CHBL's method also risks encoding the designers own biases into the neural network.

In my own experiments, I will attempt to address these issues, starting with a dataset of found images from the internet, similar to Immanuel Koh, and will show the potential to quickly assemble datasets using this method as well as addressing the possible ownership issues. My second experiment will take a contrasting approach, using a procedural design method to create custom datasets from scratch, which can allow for close control over synthesised images and be used for highly specific tasks.

I will test these datasets using a Pix2PixHD GAN, which is used for image-to-image translations, and was originally developed by NVIDIA and UC Berkeley for use in self-driving cars (Wang et al. 2018).

Input Label

Synthesised Image

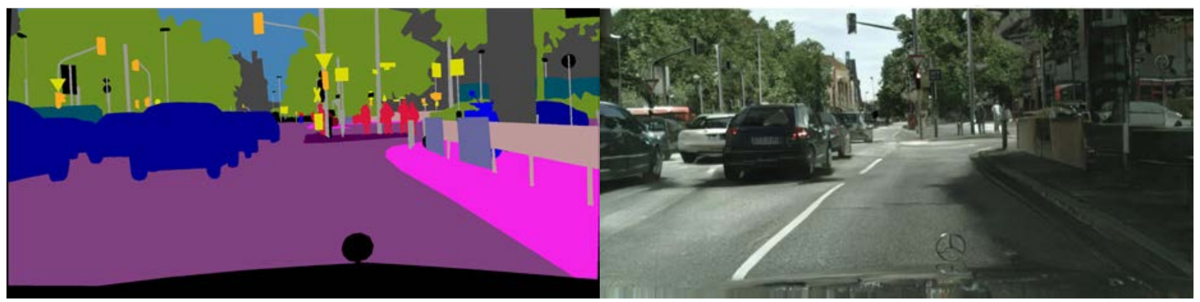

Figure 1: Pix2PixHD was developed for computer vision in self-driving cars. Wang et al. 2018.

Wang, T et al. 2018. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. NVIDIA Corporation, UC Berkeley. arXiv.

Pix2PixHD uses conditional GANs and is able to generate complex, photo-realistic outputs from simple inputs such as label maps or sketches.

My initial interest in Pix2Pix networks was spawned by DeepFaceDrawing which trained a GAN to generate realistic human faces from sketches.

The Pix2Pix GAN process is shown in figure 2. It requires a dataset of corresponding input and output images, with the generator network learning to complete the input sketch with a synthesised photo-realistic image. The discriminator then analyses whether the synthesised image is realistic enough to pass as a member of the training dataset.

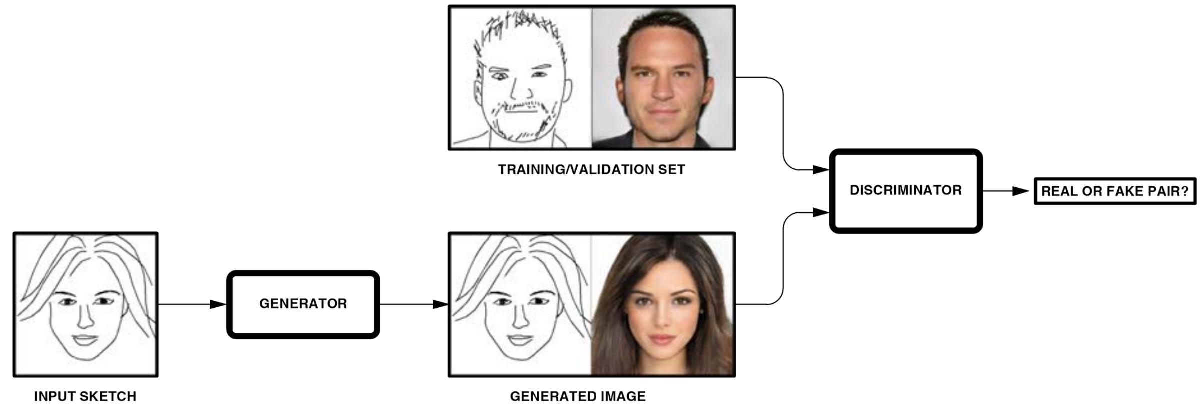

Figure 2: Pix2PixHD GAN Process. Author.

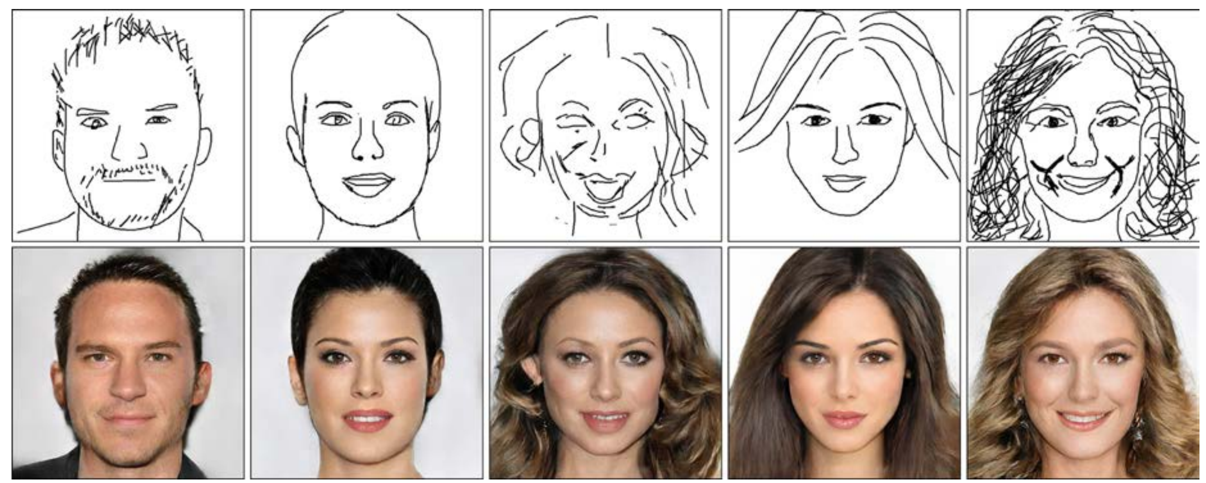

Figure 3: DeepFaceDrawing: Deep Generation of Face images from Sketches. Chen et al. 2020.

Chen, S et al. 2020. DeepFaceDrawing: Deep Generation of Face Images from Sketches

Experiment 1: Sketch a Serpentine

Dataset Collection

This experiment builds on the work of Immanuel Koh in using a GAN to design a Serpentine Pavilion. The Serpentine Pavilion offers the unique potential of having a diverse collection of architectural designs realised and photographed on the same site. This means that a coherent dataset of images of previous pavilions can be collected and used in machine learning.

The process of collecting the training images is known as 'scraping' images from the web. In this case using a Bing images search of the term, 'Serpentine Pavilion', combined with the architect's name from each year. The process was automated using the Bulk Bing Image Downloader (Ostrolucky, 2021).

It is important here to consider the ethics of this process. Web and image scraping is common practice and a fundamental tool in machine learning, however its legality and ethical standing in various circumstances is debated. There was a public outcry in 2019, when IBM released a dataset of 1 million human faces scraped from Flickr, to be used in facial recognition software. Of course the issue in this case was that people's faces were used without their permission, in order to build technologies that could be used to surveil them. There is also the question of copyright, which academics can bypass due to the non-commercial nature of their work, but professionals must be wary of (Solon, 2019). Using architectural photography, published in the public domain, issues of personal privacy are not so prevalent, but ownership and copyright issues remain. Densmore (2017) has published a list of principles of the Ethical Scraper, with the key principle, "to scrape for the purpose of creating new value from the data, not to duplicate it". Applied in the context of GANs this is important, as the resulting images from networks such as Pix2Pix are new images in their own right. In any case, in an academic setting such as this, ethics issues are minimal, but if this technology is taken into practice, these issues are likely to re-emerge. My second experiment will explore a method of designing datasets procedurally, which will remove any ethics or rights issues.

Dataset Processing

In this experiment, I use a Pix2Pix GAN with a similar process to the DeepFaceDrawing dataset, with a sketch as input and a photograph as the output. Here however, the training sketches are not drawn by hand, but using a process known as Canny Edge Detection (CED). CED uses a multi-stage algorithm to detect a wide range of edges in images and was developed by John F. Canny in 1986 (Sahir, 2019). This produces sketch-like images as shown in figure 1. As you can see, I also had isolate the pavilion from the surrounding trees and site, so that these features did not have to be sketched. This was automated using an Adobe Photoshop batch process, and then corrected manually to remove any mistakes. This manual correction was the most time-consuming part of the process, but was important to achieve accurate results. I also experimented with two variations of the canny edges. The initial edges contain no augmentation and has all of the edges present. To create the second set the image was first blurred, and the building given an outline, then was processed by the CED. This meant that there were less edges but always a

Ostrolucky, Bulk Bing Image Downloader. Github. https://github.com/ostrolucky/Bulk-Bing-Image-downloader. Accessed June 2021. | Solon, O. 2019. Facial recognition's 'dirty little secret': Millions of online photos scraped without consent. NBC News. | Densmore, J. 2017. Ethics in Web Scraping. Towards Data Science. | Sofiane, S. 2019. Canny Edge Detection Step by Step in Python — Computer Vision. Towards Data Science.

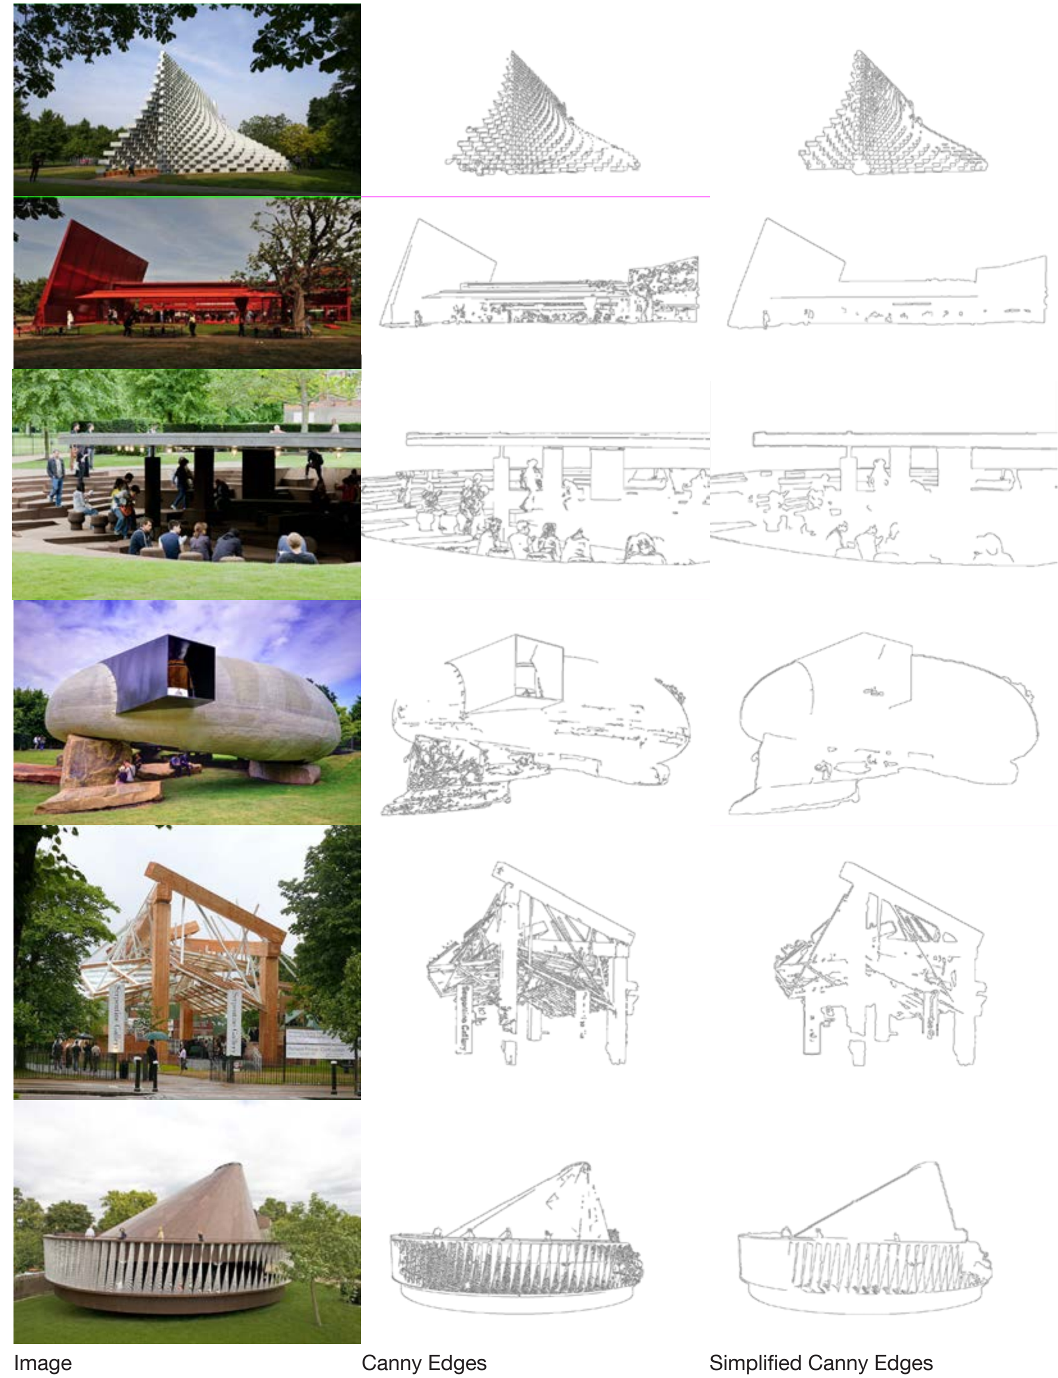

Figure 1: Training Images : 6 examples from the 350 image dataset. Author.

complete boundary, with the intention of allowing for more simple and quick sketches to be used.

Training Results

The results of the GAN training using the two variations of canny edges are shown in figures 2 and 3. Both GANs were trained for 200 epochs and use the same dataset of 354 photographs, each 1280 pixels in width. The only variation is the detail of sketch lines.

It is clear that the images synthesised using the dataset of complete canny edges were much more successful, with very accurate details. Whilst the simplified canny edges GAN is less successful, it is an interesting experiment to see whether the GAN can learn to infer the missing details in the sketch, from the remaining sketch lines. It is clear that it partially achieves this, using the correct coloration, but not quite filling in the correct textures. Even so, it could be a useful tool to test quick sketches and achieve a roughly photographic result.

Synthesised Image: Simplified Canny Edges

Synthesised Image: Canny Edges

Figure 2: Training Results: Kéré Architecture. Author.

Training Image

Simplified Canny Edges

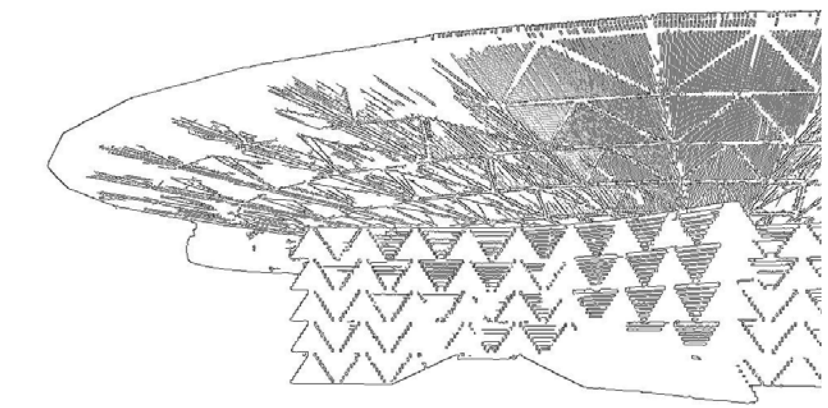

Canny Edges

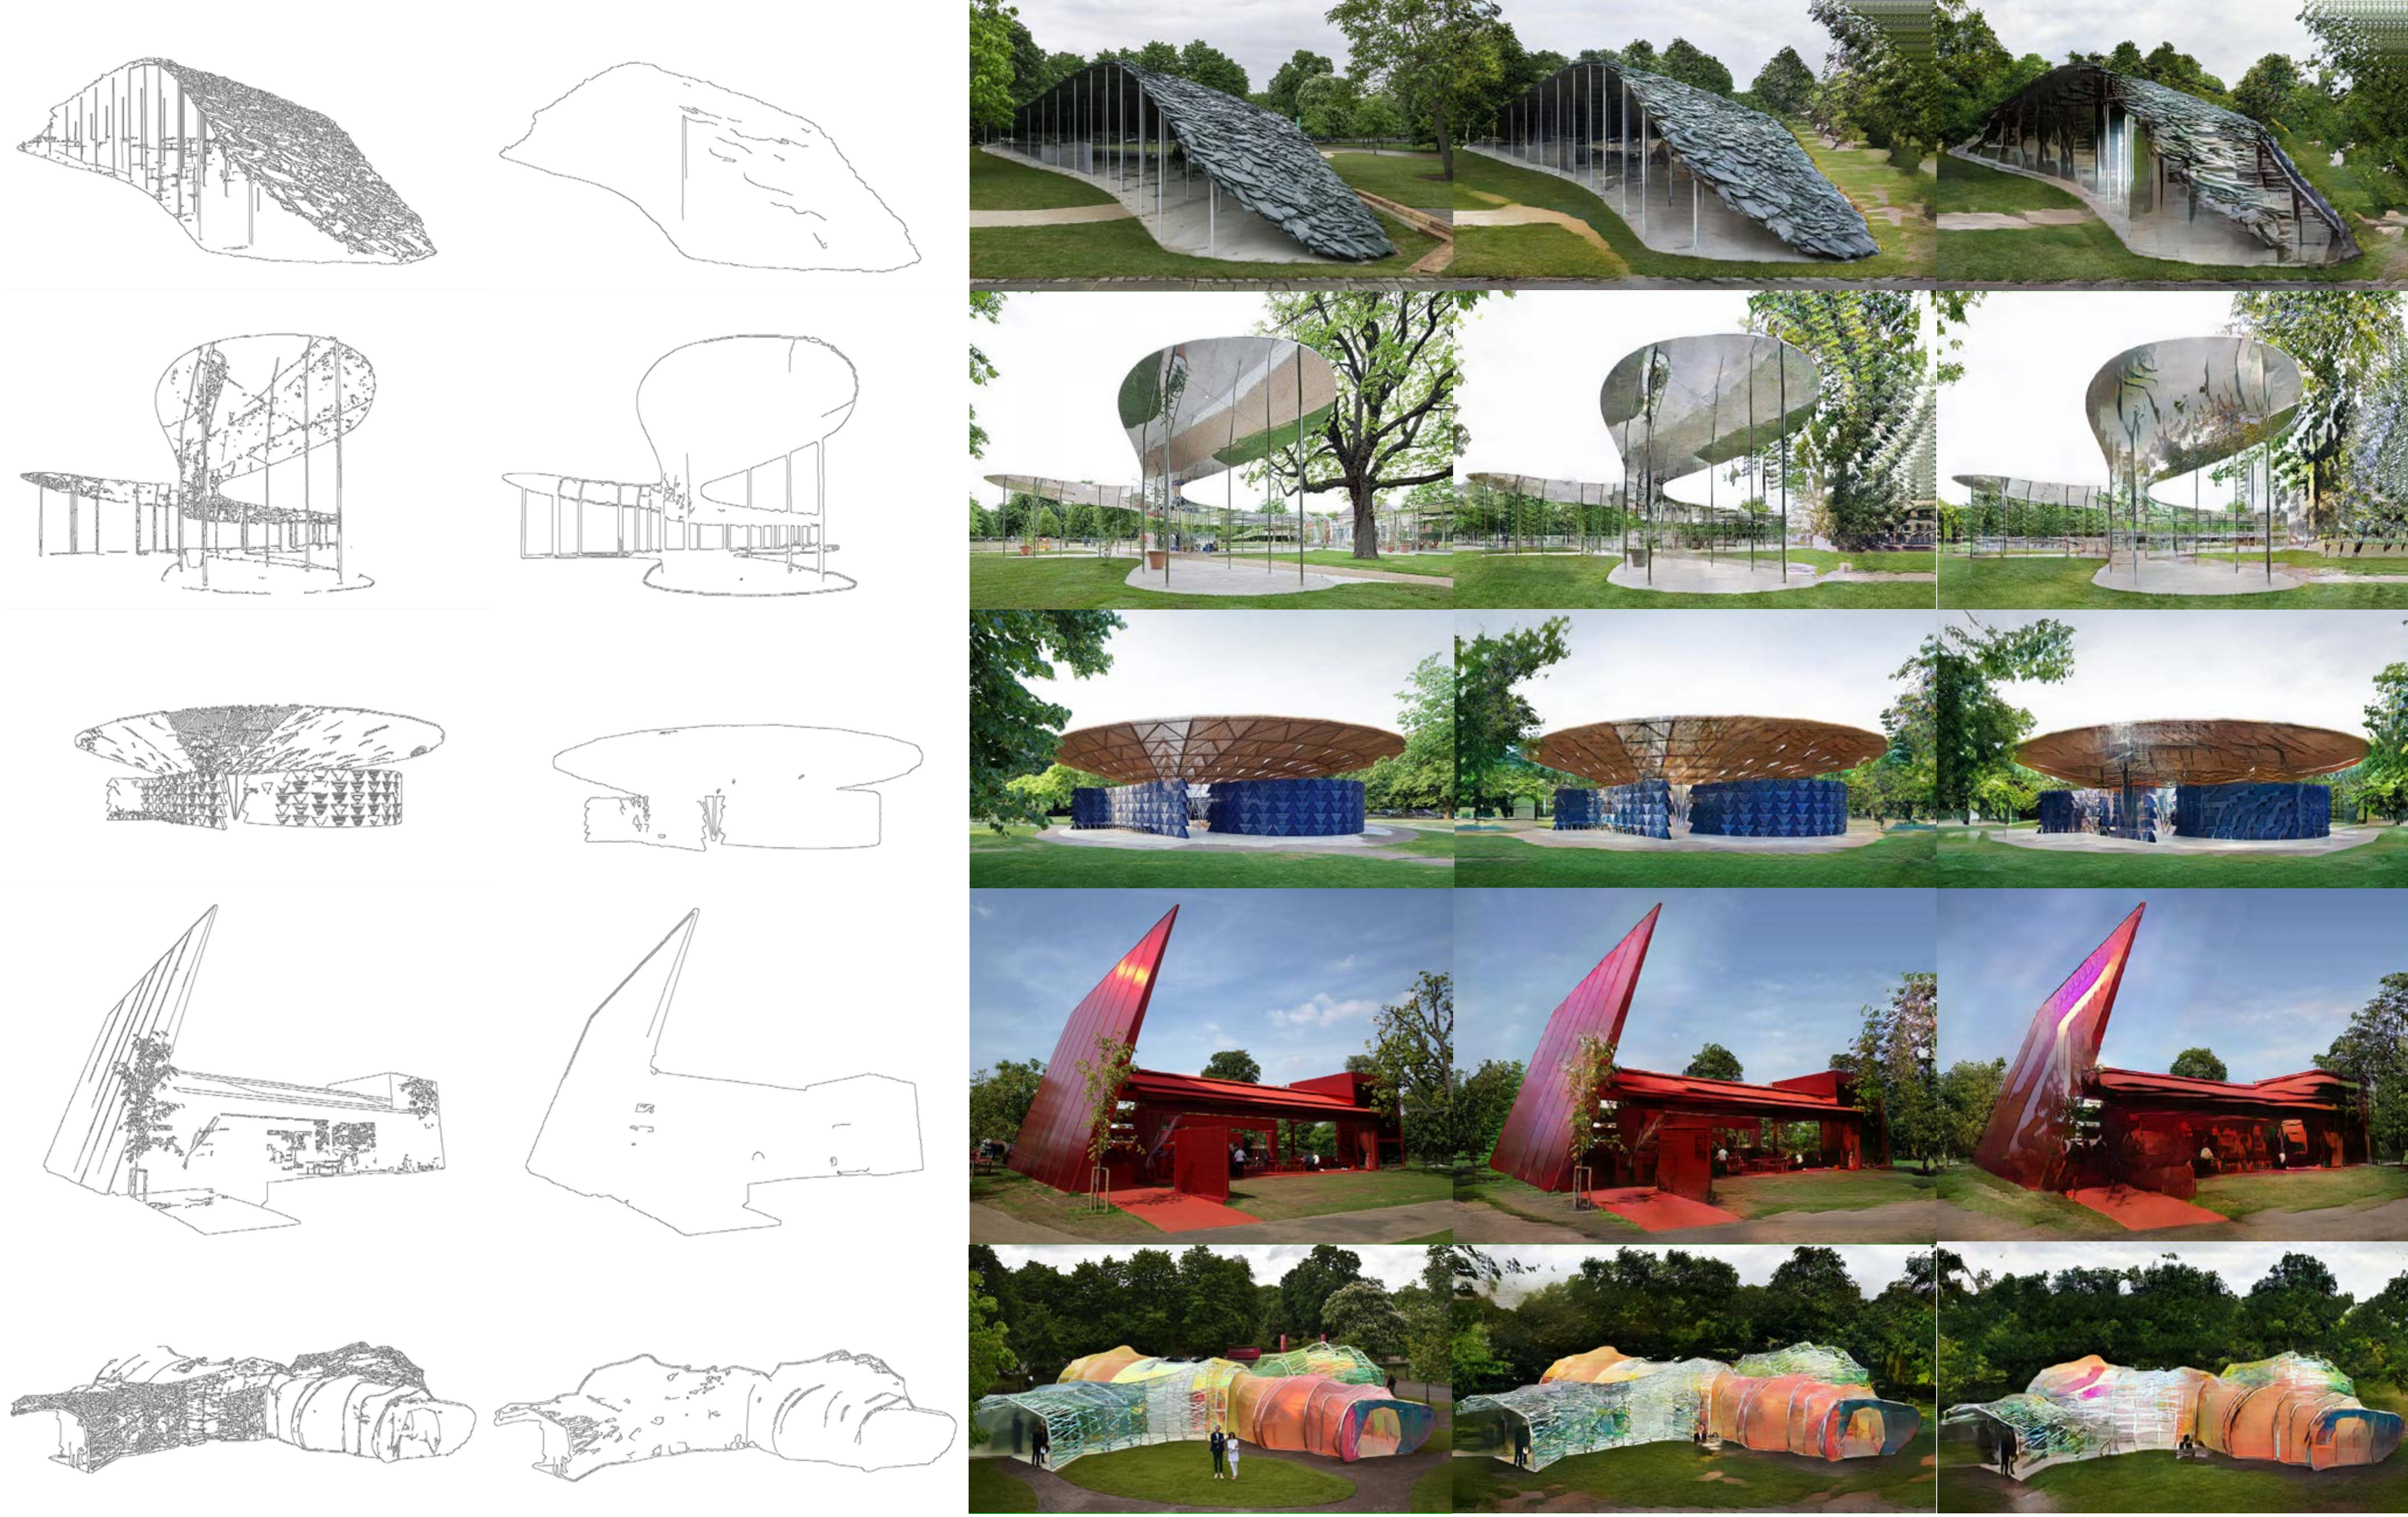

Canny Edges

Simplified Canny Edges

Training Images

Synthesised Images: Canny Edges

Synthesised Images: Simplified Canny Edges

Figure 3: Training Results: Serpentine Pavilion. Author.

Initial Testing with Custom Sketches

To test the GANs I began by drawing very simple custom sketches in Photoshop. The initial results were very distorted as in this example, although were able to be improved by adding more details to the sketch. It quickly became clear that I needed to replicate features found in the dataset more directly in order to improve the synthesised images.

In this example I tested the same sketch in both the canny edge and the simplified GAN in order to compare the results. Interestingly, unlike the training results, the quality of the canny edge GAN is not significantly better, this most likely due to the lack of detailed textures in the sketches.

Canny Edge GAN: Synthesised Image

Input Sketch

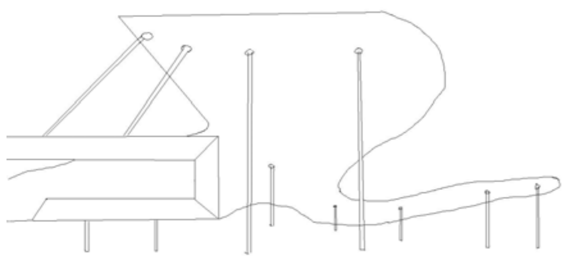

Figures 4-6 : Custom Sketch Initial Test. Author.

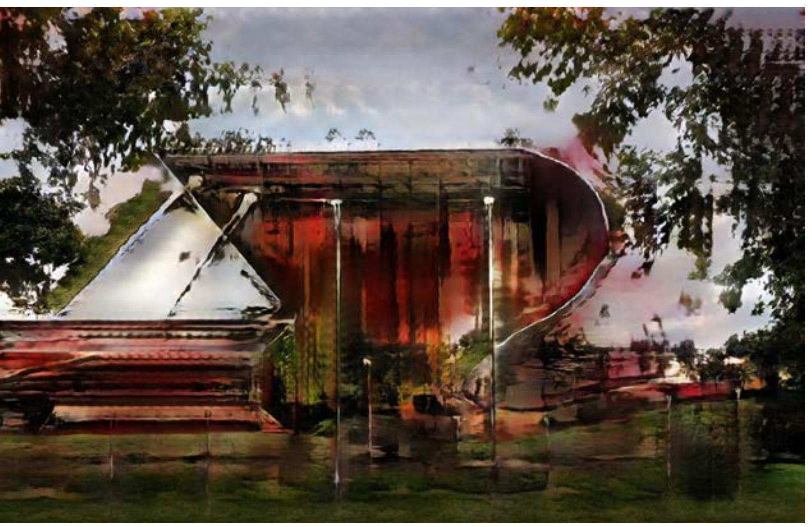

Simplified Canny Edge GAN: Synthesised Image

Figure 7 : Training Image, Jean Nouvel Pavilion. Bing Images.

With just a few small changes to the input sketch there are significant differences to the output. In the simplified GAN the synthesised image changes from the red of Jean Nouvel's pavilion to the mirror finish of the SANAA pavilion. In the canny edge GAN almost the opposite occurs, the output changing from predominantly silver to red. The simplifed GAN most closely matches my perception here, as the curving roof and dispersed columns are reminiscent of the SANAA pavilion, yet the red of Jean Nouvel's canny edge GAN likely comes from the vertical lines added to the left of the sketch.

Input Sketch

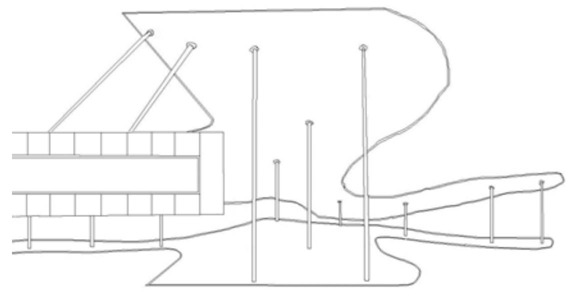

Figures 8-11 : Custom Sketch Initial Test. Author.

Canny Edge GAN: Synthesised Image

Figure 12 : Training Image, SANAA Pavilion. Bing Images.

Simplified Canny Edge GAN: Synthesised Image

Hybrid Sketches

From testing my early sketches I found that it was possible to create the most successful results when I reused textures and forms present in the original pavilions, creating a series of hybrid results. The GAN applies colours and textures reminiscent of the training data when it recognises certain forms and patterns.

These examples were all generated by the simplified GAN, using minimal details in the sketches. Elements can be recognised from pavilions in training set, yet are abstracted and blended with other influences.

Figure 13 : Training Images. Bing Images.

Figure 14 : Synthesised Image

Figure 15 : Sketch Input

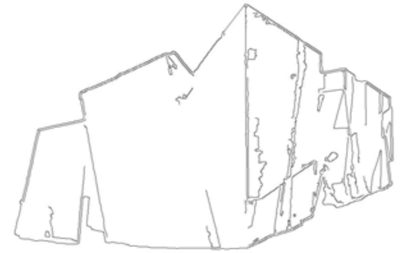

Figure 16 : Sketch Input

Figure 17 : Synthesised Image

Figure 18: Sketch Input

Figure 19: Synthesised Image

Developed Designs

These developed works build on the previous hybrid sketches, this time using the canny edge GAN. The speed at which the GAN can synthesise an image gives almost instant feedback when sketching, allowing for the sketch to be incrementally improved based on the output image.

Once again, these examples re-purpose features present in previous pavilions, for example Figure 22 uses the colours and textures of the SelgasCano pavilion whilst figure 23 combines a cantilevered roof with the texture of Kéré Architecture's pavilion, with the sunken seating area of Herzog and de Meuron and Ai Weiwei's pavilion.

Figure 20: Training Images, SelgasCano. Bing Images.

Figure 21: Sketch input. Author.

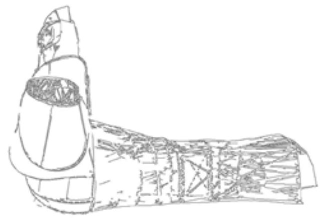

Figure 22: Synthesised Image. Author.

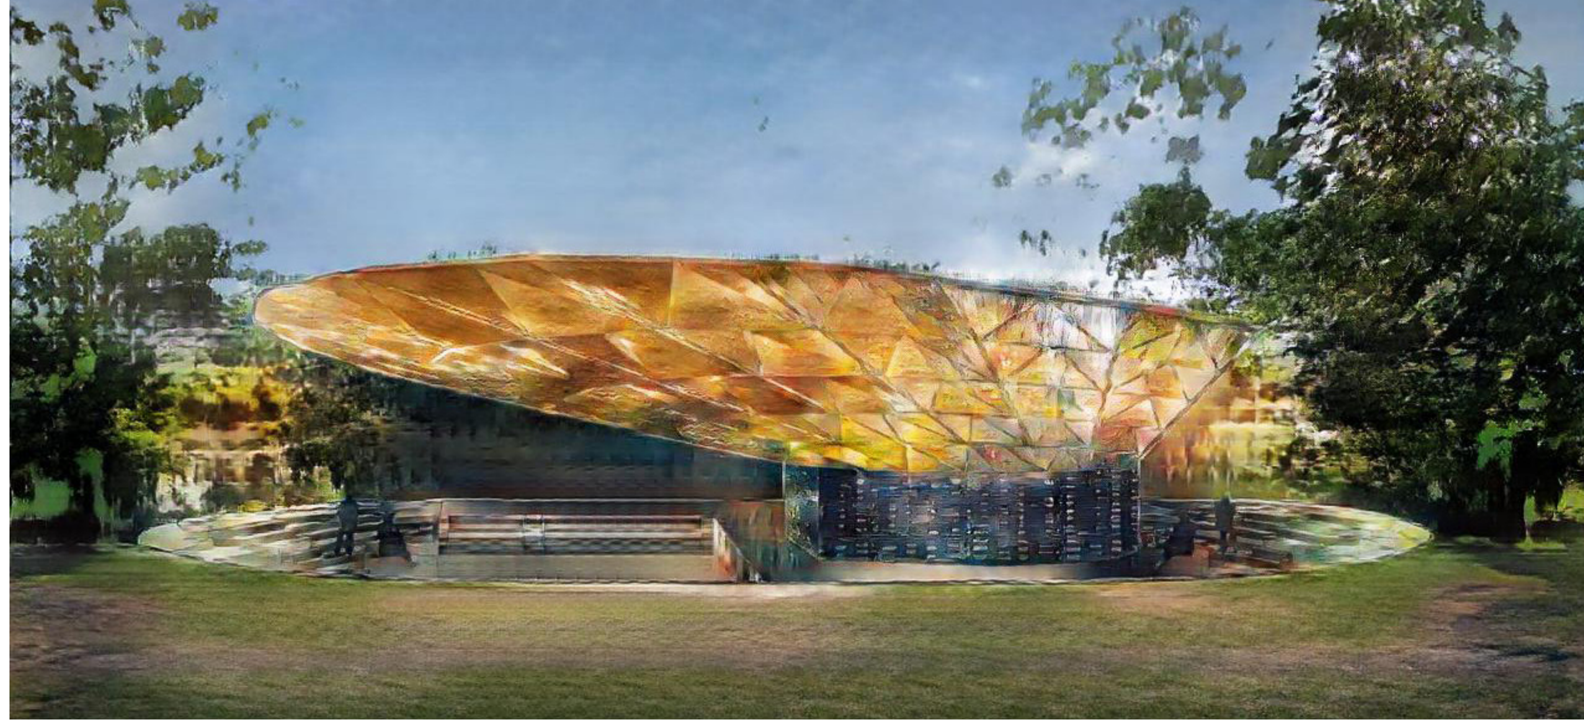

Figure 23: Developed Synthesised Image. Author.

Figure 24: Training Images, Kéré Architecture, Herzog and de Meuron. Bing Images.

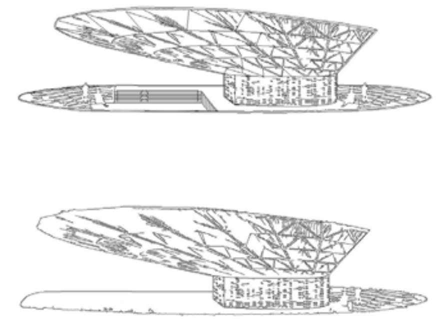

Figure 25: Initial and Developed Sketch inputs. Author.

Figure 26: Initial Synthesised Image. Author.

Developed Designs

This final design contains significant detail built up over a number of sketch iterations. It borrows the louvres from Olafur Eliasson's 2007 pavilion, however the resulting image is red, suggesting the overall form has been recognised as Jean Nouvel. It also contains a roof that uses the texture of Junya Ishigami's slate pavilion.

Figure 27: Training Images, Olafur Eliasson, Junya Ishigami, Jean Nouvel. Bing Images.

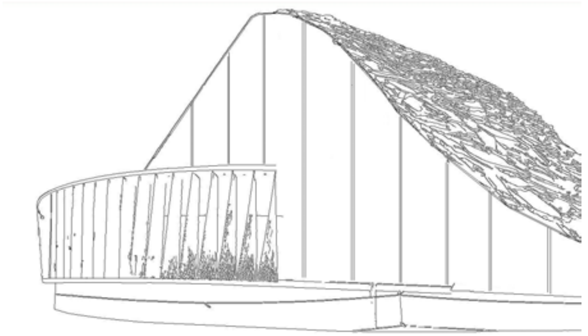

Figure 28: Final Sketch Input. Author.

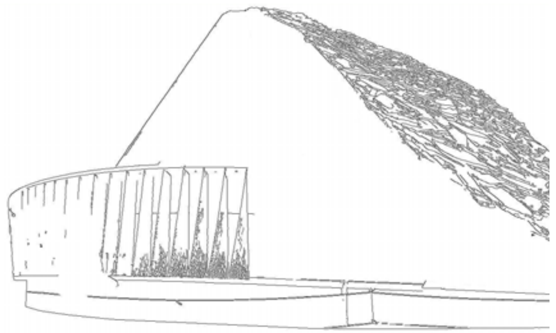

Figure 29: Initial Sketch Input. Author.

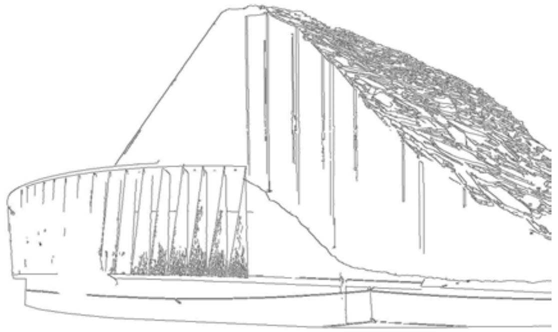

Figure 30: Developed Sketch Input. Author.

Figure 31: Final Synthesised Image. Author.

Figure 32: Initial Synthesised Image. Author.

Figure 33: Developed Synthesised Image. Author.

Testing Perception

This final experiment tests the GANs perception of well-known architectural landmarks. If Morphosis or Le Corbusier were to design a Serpentine Pavilion what would the result be? Which past pavilions would their works be perceived as? In this case the purist white Villa Savoye is translated into the dark timber-clad finish of Olafur Eliasson's pavilion, whilst the bold, stone-clad forms of Morphosis's Perot Museum are finished in the metallic red of Jean Nouvel.

To understand the reason for this interpretation by the GAN we can analyse the similarities in form between these landmarks and their corresponding pavilions (see figure 27). Connections that we might not have otherwise made become clear, for example the elevated ramp of Eliasson's pavilion corresponding to the strip window of Villa Savoye, whilst the sharp form of Nouvel's pavilion corresponds quite clearly with that of the Perot Museum.

Figure 35: Synthesised Villa Savoye. Author.

Figure 34: Villa Savoye. Bing Images.

Figure 37: Synthesised Perot Museum. Author.

Figure 36: Perot Museum. Bing Images.

Analysis and Application

The Sketch a Serpentine experiment has addressed each of the three thesis questions through its preparation of datasets, and its exploration of a neural network's perception and creation. It has shown how datasets can be built and augmented using images scraped from the web and explored the issues regarding privacy and copyrights that can arise from this. The use of a Pix2Pix GAN has shown how the age-old process of sketching can be augmented to provide near live feedback, in the form photographic renderings, providing a new creative stimulus to the design process. As you sketch, the synthesised images are constantly challenging your perception: unexpected textures or finishes can result in new design ideas, whilst inaccuracies or distortions could spur new avenues of exploration.

The GAN has shown that it can learn to associate certain sketched forms and textures with images from the training set, but it has become clear that the small size of the dataset (354 image pairs) has limited the successful outputs. Only the select features and details that exist in the training set can be replicated, which has limited the sketching possibilities significantly. With a larger and more diverse dataset the GAN could learn to interpolate many more features. If on the scale of IBM's facial recognition dataset of 1 million images, the possibilities would be far-reaching. A dataset of 1 million buildings would be diverse enough to allow for almost any well-defined sketch to be recognised and translated into a photorealistic image. This suggests that a GAN could be used to hugely speed up the current process of architectural rendering. A digital sketch could be processed into a photographic image in near real time. Equally, a digital model could be exported as a wire-frame image, or similar coded label map, which could then be translated by the GAN into a photorealistic rendering. This would bypass the time-consuming and resource-heavy rendering processes available today. There would need to be more control over the final output and an ability to select materials and finishes, but this could be accomplished by replacing the sketch input with a coded label map. Creating these kind of label maps will be explored in the next design experiment.

Experiment 2: Procedural Datasets

This series of experiments continues the exploration of how architects can best prepare datasets for use with machine learning, taking a contrasting approach to the web scraping of the previous Serpentine dataset. Rather than using found images, I will be designing the datasets myself. This can overcome a number of issues, firstly removing any ethical dilemmas surrounding data privacy and ownership. More fundamentally, it removes any reliance on past or completed works and so offers the space for more speculative works. Of course, it also brings its own challenges. Deep Learning requires large datasets of at least hundreds or thousands of images, which can be a daunting prospect when each must be designed from scratch. Procedural design offers the potential to overcome this challenge. Rather than designing a single entity, using procedural design you are able to define a 'recipe' that can be adjusted, refined and repeated to create similar, yet unique results. This can be described as a procedural, 'species', of designs which can adapt to various contexts and inputs.

Specifically, using a system known as Procedural Dependency Graphs (PDG) it is possible to distribute the generation and rendering of each variation of the procedural species, allowing for the quick generation of large datasets.

Using this method I have created three datasets that have been tested by training a Pix2PixHD GAN. They all explore methods of controlling the output image using various types of input, beginning with a sketch, and developing into the use of colour coded input labels. The aim is to maximise the control over the output image, creating an augmented design process, with constant feedback between the designer and the neural network.

Dataset 1

The first dataset consists of 1024 procedural objects and has been tested in a similar way to the Serpentine dataset, using canny edges as the input and a rendered image as the output.

The training results show that the GAN is able to recreate the training images, albeit with some distortions (figure 2). The areas without any canny textures are most difficult for the GAN to interpolate, and so some gaps in the rendered image occur.

When tested with hand drawn sketches, as in figure 3, the synthesised images successfully recreate the overall form, yet there are many gaps and distortions.

The problem here is that the micro textures of the procedural objects have very specific patterns that create intricate canny edges. These edges are difficult to replicate when sketching, and so the resulting images lack the detail of the original objects.

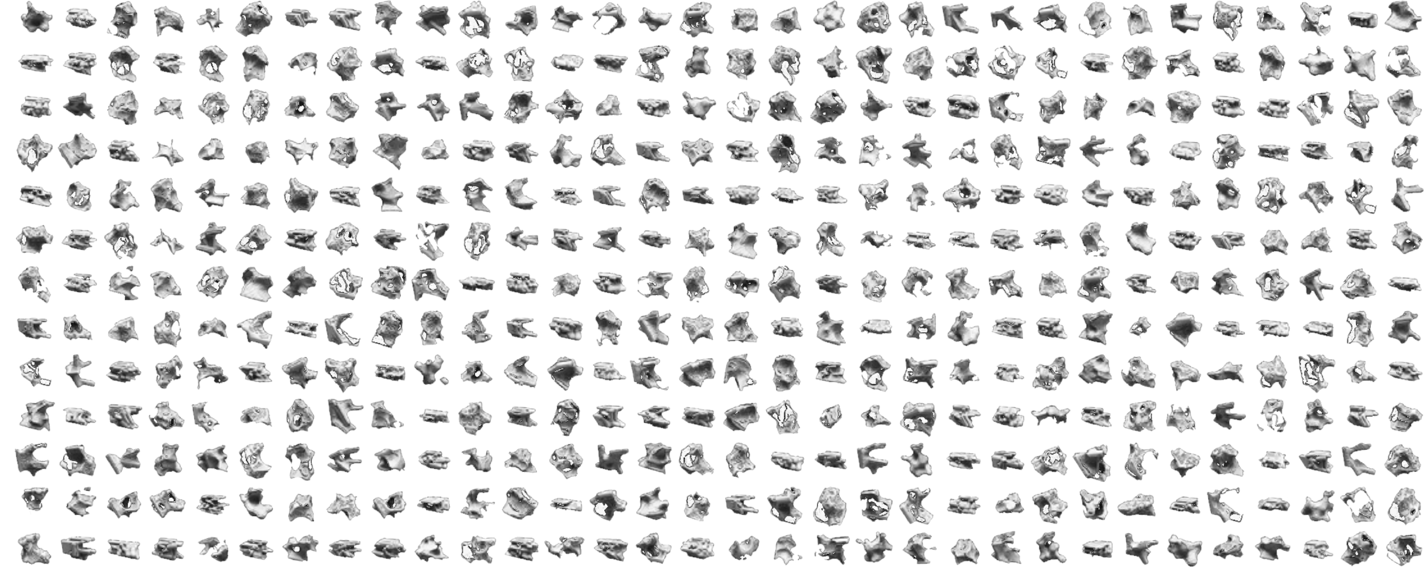

Figure 1: 416 of 1024 Image Training Dataset. Author.

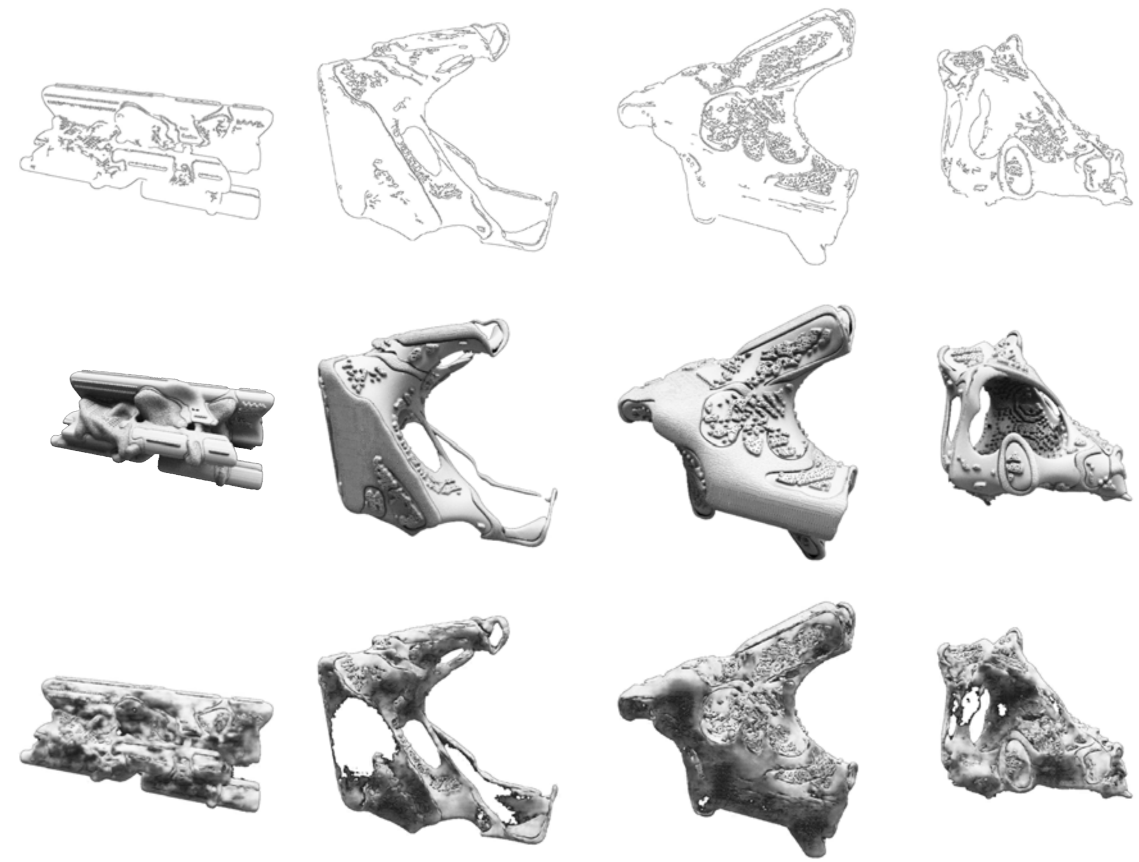

Figure 2: Pix2PixHD GAN training results. Author.

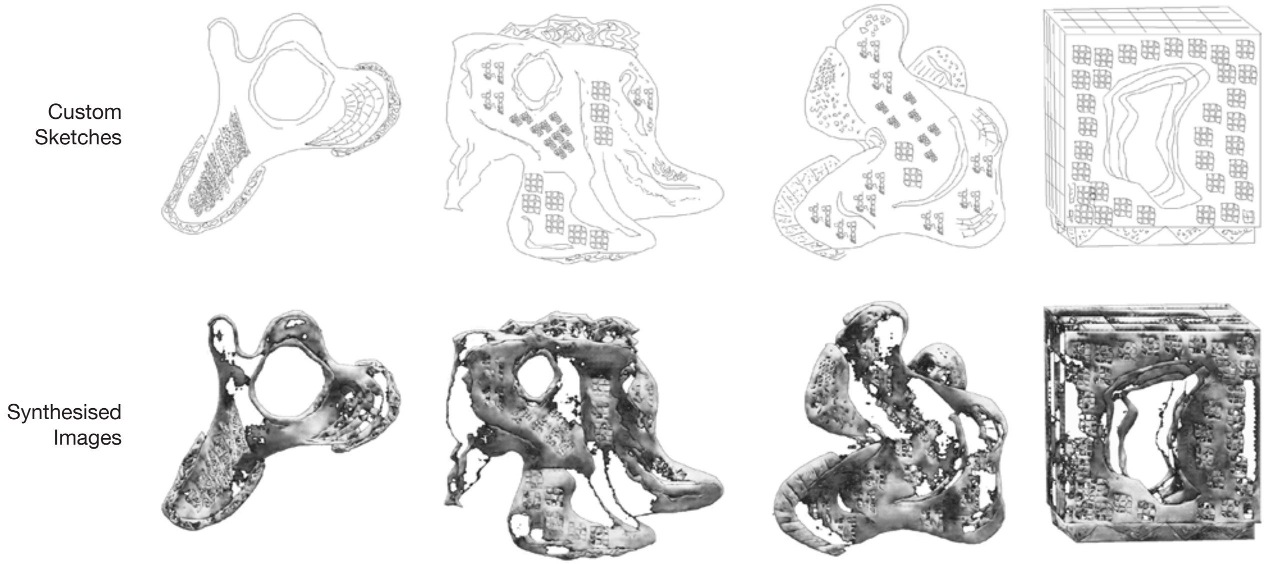

Figure 3: Pix2PixHD GAN testing with custom sketches. Author.

Dataset 2

The second dataset builds on the first by adding colour coding to the sketch input and sketching in plan view, rather than perspective. The idea is to use the sketch input to generate a 3D massing model.

In this case the sketch inputs are not generated using canny edges, but procedurally. The training models are sliced to generate lines similar to contours, which are then colour coded, to represent heights. In this case, the output image is no longer a rendered image, but a greyscale depth map, which can be used to generate a massing model. In this way, the designer can sketch in 2D whilst receiving live feedback in the form of a 3D model.

Given such a specific task of translating between coloured contours and a gridded depth map, the GAN is able to very accurately synthesise new depth maps as shown in figure 5.

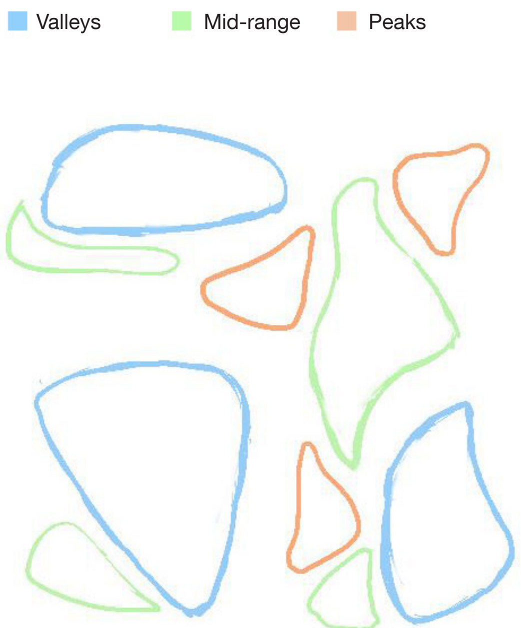

Figure 4: Custom input sketch, colour coded. Author.

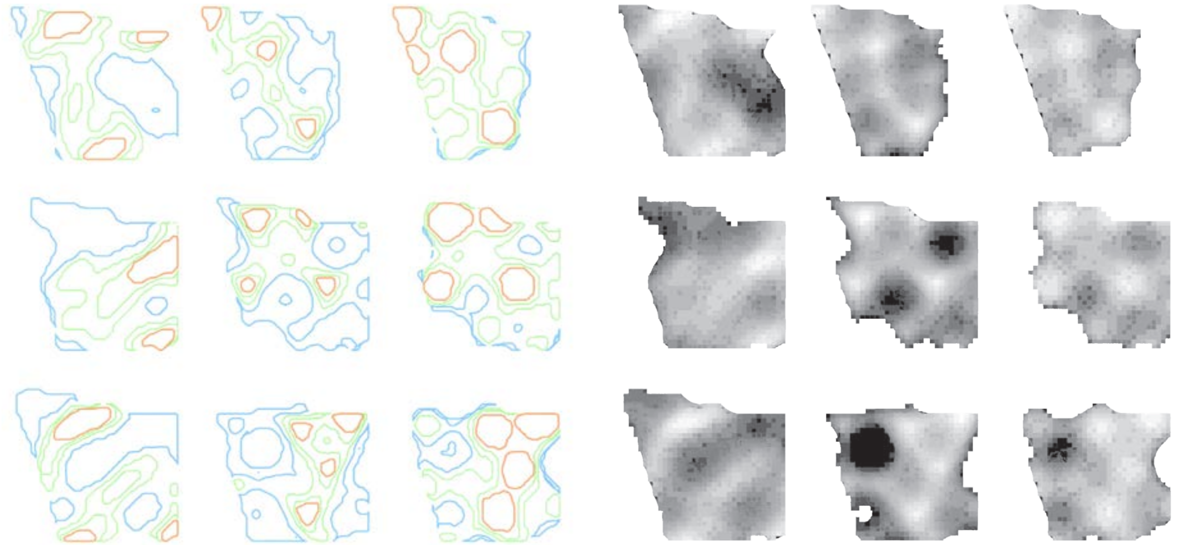

Figure 7: Training Set of coloured sketches and depth maps. Author.

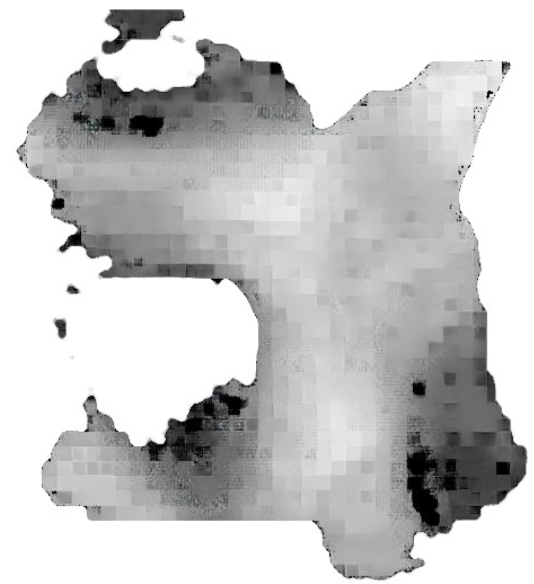

Figure 5: Synthesised greyscale depth map. Author.

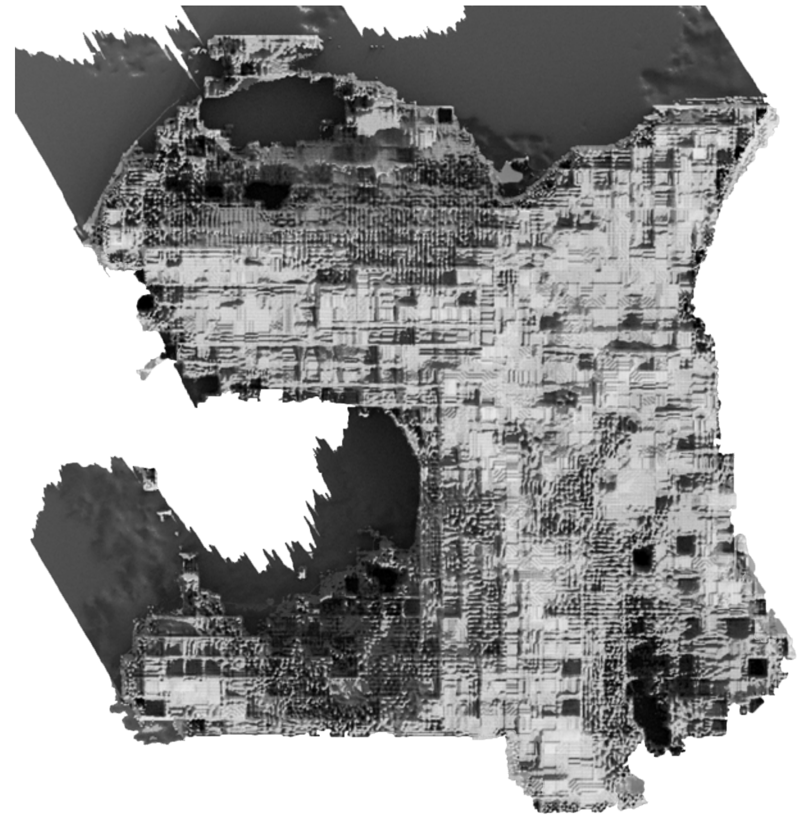

Figure 6: Massing model from synthesised depth map. Author.

Dataset 3

The third dataset takes the idea of using colour coding further, using a label map instead of a sketch. It questions whether the GAN can be used to generate floor plans, based on a series of coloured labels that identify various types of spaces. For example, the lightest areas of the map are day-lit areas, which are intended to be open plan, live or work spaces. Dark grey areas receive least light, and should contain more pocketed, service spaces. Orange represents the areas which receive the most direct sunlight, and so these are given over to the non-human.

The test is to see whether the GAN can learn to generate the correct density of walls and open spaces, based on the input colours. It also must learn to replicate the correct colours and textures present in the training plans.

As seen in figure 10, the GAN has successfully learnt the meaning of each colour, and has interpolated the location of the walls, synthesising a seemingly feasible floor plan that closely resembles the training set. Perhaps the weakest area is in the day-lit spaces, where too many of the walls are disconnected from each other, creating open plan spaces that appear too large.

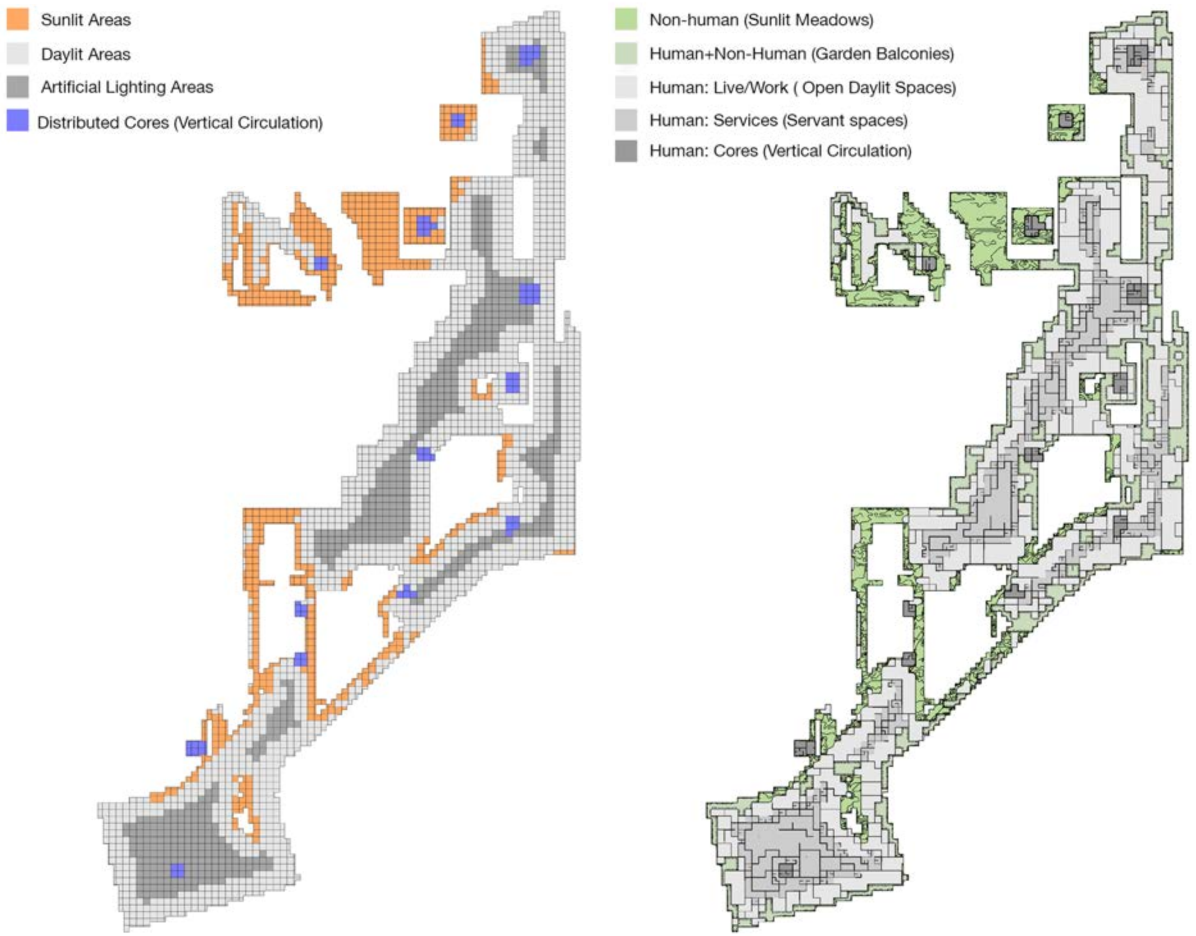

Figure 7: Training Set Example: Label Map Input, Floor Plan output. Author.

Figure 8: Training Set Examples: Label Map Input, Floor Plan output. Author.

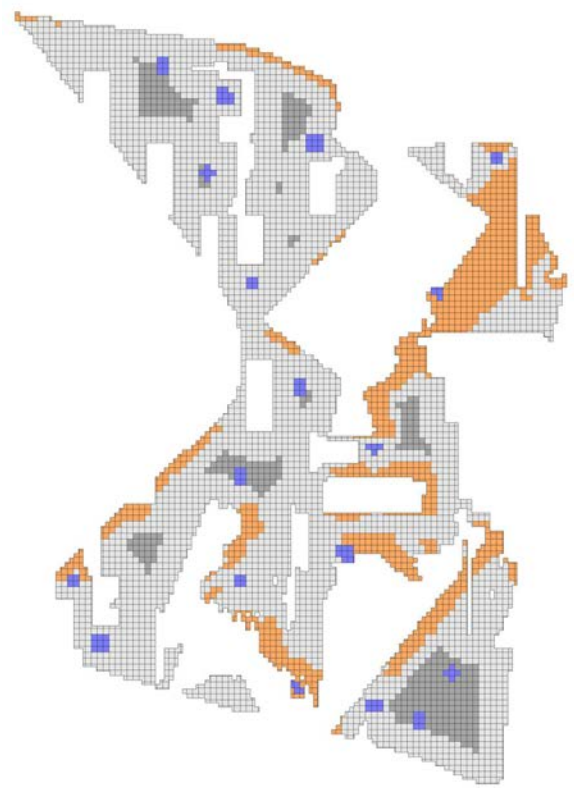

Figure 9: Testing: Label Map Input. Author.

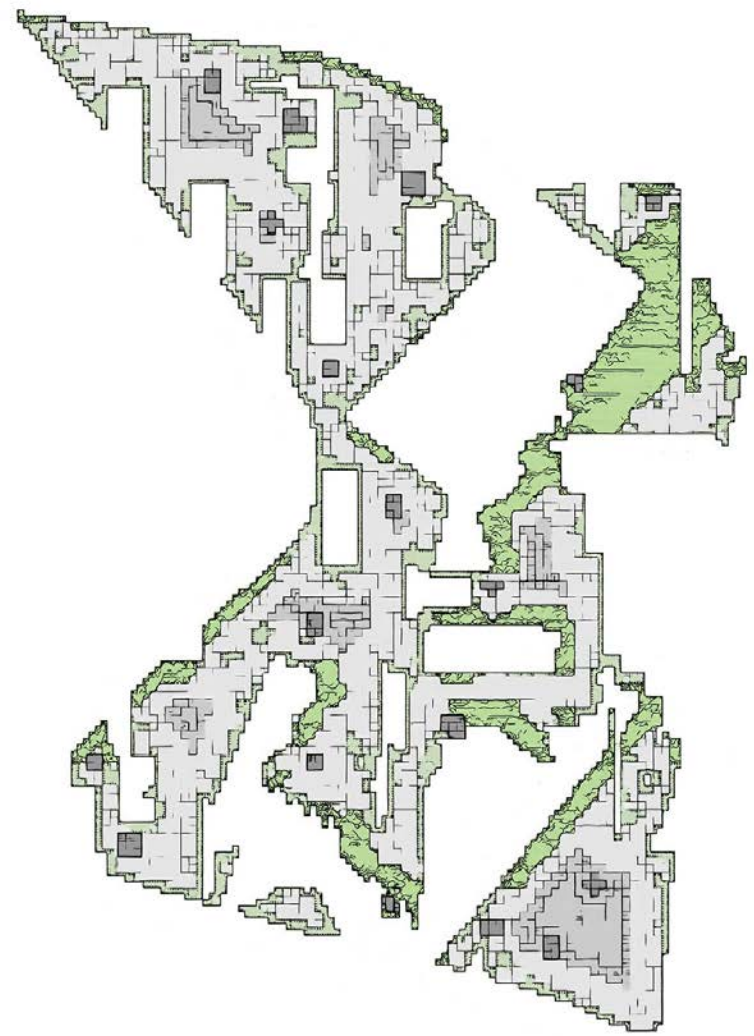

Figure 10: Testing: Synthesised Floor Plan. Author.

Analysis

These experiments have shown that procedural datasets can be a powerful tool for developing design tools using a Pix2PixHD GAN. Procedural methods allow for the creation of large datasets of images with relative ease, without having to rely on data sourced from the web.

By encoding information into the input images using colours, the designer can achieve close control over the synthesised images. This can allow for an augmented design process in which the designer can iterate and improve his sketch, or input map, and receive live feedback.

This type of dataset seems to be well-suited to very specific tasks, such as the generation of a massing model or floor plan. It is a very different use of GANs to the previous experiment of the Serpentine Pavilion, and the work of Immanuel Koh and CHBL. This comes from the specificity of the datasets used; whilst each variation of a procedural species is different, they still have very similar characteristics. The differences between two procedural plan drawings are much less significant than between two Serpentine Pavilions. This means that the synthesised images are much more predictable, and could be used to create very reliable design tools. Yet, what is lost in these tools is the perceptual challenge which earlier datasets offered. When sketching a plan using Procedural Dataset 3, you are not going to achieve a significantly unexpected result, and without this element of surprise, one of the key components of creativity is removed. This does not necessarily negate the usefulness of this dataset as a design tool, but the creative stimulus of the design process must be found elsewhere.

Conclusion

AI was devised as a technology to replicate the human brain, and since its inception was put in competition with human intelligence. Throughout its portrayal in popular fiction and early testing, AI took on the role of the antagonist of the human, always threatening to assert its superiority through dystopian confrontations or hyped up chess matches. Yet as AI continued to be developed, and the Turing test repeatedly failed, it became clear that the intelligence of machines does not have the same broad intelligence as the human brain. Instead, AI can excel to superhuman levels of intelligence, but only in highly specific statistical tasks.

Deep learning neural networks are the most advanced form of AI, and have the ability to build their own perception though the analysis and interpretation of huge quantities of data. Just like our own human perception, machine perception is completely dependant on its past experience, and can inherit biases and prejudices from its training data. At the same time, it can also be used as a method of challenging our own perception, allowing us to see from a perspective outside of our own, free from our personal or cultural preconceptions. This kind of human-machine experimentation can provide a stimulus for creativity, allowing a designer to produce unexpected and surprising results. The creativity of this process comes from the collaboration between human and machine, using the statistical power of AI in combination with the human capacity for contextualisation and assessment.

The proceeding case studies have provided examples of how machine perception can be used to augment the creative process. Immanuel Koh demonstrated how neural networks build a perception. He showed that they can interpolate and distinguish between the work of specific architects, as well as iterate novel design possibilities. Coop Himmelb(l)au has used neural networks to further their investigations into unconscious design processes. Neural networks are quite literally unconscious designers, interpolating features from large datasets, without awareness of their actions. They are a technology and a tool not just to be used by humans, but to be interacted with in a constant feedback loop.

My own experiments have investigated the design of a broad range of datasets, and have tested their application using a Pix2PixHD GAN. Sketch a Serpentine built on the work of Immanuel Koh and explored the curation, and augmentation, of datasets from the web, as well as the associated ethical questions. It exhibited how the traditional design process of sketching could be augmented, using AI to spur new design avenues through its varied, and sometimes distorted, perception of sketch lines. It also speculated on future rendering tools that could be built using this technology, and the scale of datasets that would be needed.

The procedural dataset experiments explored the creation of custom datasets, and showed how sketch inputs could be developed into label maps, encoding more information and allowing for very close control over synthesised images. In this scenario Pix2Pix could be used to build very reliable and specific tools, but lost some of its freedom and potential for creative accidents and unexpected outcomes. This is perhaps one drawback of using Pix2Pix GANs versus the latent space model used in the latest phase of DeepHimmelblau. When exploring the latent space, no specific input is required, allowing the designer to take a hands off approach, essentially becoming a curator rather than creator, and giving over more agency to the neural network.

In any case, these differences show the versatility of neural networks and specifically GANs, and the variety of ways that they can be used. It is clear that Pix2PixHD could be extremely useful and proficient in specific tasks, such as generating photo-realistic renderings from sketches, or floor plans from label maps. At the same time, GANs such as the CycleGAN, used by Deep Himmelblau, are better suited to be used for creative stimulation, as they can create truly unexpected and surprising results. It is these accidental moments, that in our perception could be considered mistakes, or hallucinations, that can often have the most value in sparking new design ideas or avenues to explore.

It is clear that the future potential of AI in architectural design lies in the close collaboration and feedback loop between human and machine. In this way the relative strengths of each can amplify and challenge each other; the intuitive and broad intelligence of the human can be extended to super-human levels in particular tasks, using the analytical power of data analysis that AI offers. Machine perception can be used to challenge our biases and synthesise novel possibilities, whilst human judgement can set the agenda and asses the value of results. The full effects of this collaboration are not yet certain, but may drive the continuing evolution of the human, toward Fuller's cyborg vision of a unity between man and apparatus.

Bibliography

Akst, D. 2020, AI in Popular Culture: How Much Do You Remember?. https://www.wsj.com/articles/ai-in-popular-culture-how-much-do-you-remember-11604437200.

Bolojan, D. 2020. Meet Deep Himmelb(l)au. https://nonstandardstudio.com/2020/01/26/meet-deephimmelblau/.

Bolojan, D. 2021. FIU DDES Lectures_Theories of the Digital_Session4: AI and Creativity. Youtube.

Bolojan, D. 2021. FIU DDES Lectures: Theories of the Digital Session 7_ AI & The Architectural Office of the Future. Youtube.Main text here.

Brownlee, J. 2020. How to Explore the GAN Latent Space When Generating Faces. Machine Learning Mastery.

Campo, M., and Manninger, S. 2019. "Imaginary Plans." Proceedings of the 2019 ACADIA Conference - Ubiquity and Autonomy.

Chaillou, S. 2019. 'The Advent of Architectural AI'. Harvard Graduate School of Design.

Chen, S et al. 2020. DeepFaceDrawing: Deep Generation of Face Images from Sketches

Culkin, J. 1967. A Schoolman's Guide to Marshall McLuhan. Saturday Review.

Densmore, J. 2017. Ethics in Web Scraping. Towards Data Science.

Du Sautoy, M. 2019. The Creativity Code: Art and Innovation in the Age of AI. UK: Harvard University Press.

Flusser, Vilém. 2000. Towards a Philosophy of Photography. London: Reaktion Books.

Hoteit, A. 2015. Deconstructivism: Translation From Philosophy to Architecture. Canadian Social Science, 11(7), 117-129.

Kasparov, G. 2018. 'Chess, a Drosophila of Reasoning'. Science 362, no. 6419.

Koh, I. 2020. Serpentine2020?. https://artificial-architecture.ai/?p=140.

Koh, I. 2021. Serpentine2020?. FIU DDES Lectures: Theories of the Digital Session 6 on AI & Architectural Design. https://www.youtube.com/watch?v=u-8KO_7ycgE

Kurzweil, R. 2005. The Singularity is Near: When Humans Transcend Biology. New York: the Penguin Group.

Kuszewski, A. 2009. The Genetics of Creativity: A Serendipitous Assemblage of Madness. Metodo.

Memo, A. et al. 2019. Learning to see: you are what you see. SIGGRAPH 2019. Association for Computing Machinery, New York, Article 13, 1–6

Metz, C. 2016. In Two Moves, AlphaGo and Lee Sedol Redefined the Future. Wired. https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/.

Mitchell, M. 2019. Artificial Intelligence: A Guide for Thinking Humans. Penguin UK.

O'Neil, C. 2017. The Era of Blind Faith in Big Data Must End. TED Talks. https://www.ted.com/talks/cathy_o_neil_the_era_of_blind_faith_in_big_data_must_end/.

Ostrolucky. 2021. Bulk Bing Image Downloader. Github. https://github.com/ostrolucky/Bulk-Bing-Image-downloader. Accessed June 2021.

Russell, S., Peter N. 2010. Artificial Intelligence: A Modern Approach. 3rd ed. Prentice Hall.

Seth, Anil. 2017. "From Unconscious Inference to the Beholder's Share: Predictive Perception and Human Experience." PsyArXiv.

Sofiane, S. 2019. Canny Edge Detection Step by Step in Python — Computer Vision. Towards Data Science.

Solon, O. 2019. Facial recognition's 'dirty little secret': Millions of online photos scraped without consent. NBC News.

Silver, D, Hassabis, D. 2016 AlphaGo: Mastering the ancient game of Go with Machine Learning. Google DeepMind. https://ai.googleblog.com/2016/01/alphago-mastering-ancient-game-of-go.html.

Silver, D, et al. 2018. AlphaZero: Shedding new light on chess, shogi, and Go. https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go.

Sternberg R. J. 2012. Intelligence. Dialogues in clinical neuroscience, 14(1), 19–27.

Stezano, M. 2017. In 1950, Alan Turing Created a Chess Computer Program That Prefigured A.I. https://www.history.com/news/in-1950-alan-turing-created-a-chess-computer-program-that-prefigured-a-i

Tattersall, I. 2020. "Homo sapiens." Encyclopedia Britannica.

Tegmark, Max. 2017. Life 3.0: being human in the age of artificial intelligence.

Turing, A. M. 1950. Computing machinery and intelligence. Mind, 59, 433–460.

Wang, T et al. 2018. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. NVIDIA Corporation, UC Berkeley. arXiv.

Zylinska, J. 2020. AI Art: Machine Visions and Warped Dreams. London: Open Humanities Press.