.Hyperrealism – The Influence of Changing Technologies on Architectural Rendering

by Lee Jackson

Master of Creative Technologies

Supervisor: Stefan Marks & Andrew Burgess

Date Submitted: February 16, 2023

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly define in the acknowledgements), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Unless otherwise stated, all images in this document were created by me. Images that were sourced were done with either the express permission of the original author or fall under fair use.

Lee Jackson

Acknowledgements

I would like to thank my wife Shannon, for encouraging me to do this thesis and being a source of inspiration as she completed her PhD.

I would like to thank my supervisors, Stefan Marks and Andrew Burgess for constantly pointing me in the right directions and steering me clear of the wrong ones.

Abstract

Hyperrealism is a concept that has evolved from paintings depicting real life portraits to digitally created content. Various industries are employing hyper-realistic techniques, however, the focus of this research will be on the architecture and architectural rendering fields due to the work being created. With an increase of hyperreal advertisements coming from these fields, enabled by the rapid development of computer graphic technology and hardware that is used to produce these images, there has to date been very little scholarly research into how hyperreal images are being produced and the effects that these images can have on consumers.

The ability to create hyper-realistic images can deceive and/or create unrealistic expectations, and some renders are even starting to breach various marketing and advertising laws through false advertising. Are architecture firms aware of these factors that can be produced with hyperreal renders or are they just a by-product of the advances that have happened with the technology used to produce their renders?

This thesis explores the history, concept and definition of hyperrealism and the potential effects that hyperreal imagery can have. This thesis also describes the steps necessary to produce hyper-realistic architectural images to inform future students, artists, and creatives looking to create photoreal and hyperreal environments. This thesis will also look at existing architecture firms in New Zealand and examine the images they are using to promote their practices online to determine if they contain hyperreal elements and how the respective companies are promoting their images.

Table of Contents

List of Figures

Chapter 1. Introduction

Realism in the architectural rendering community is becoming much more commonplace as technology gets better and as 3D artists are becoming more skilled. What problems are created when rendering surpasses the real and starts to venture towards hyperrealism? What effects, if any, do these perceived problems have on the professionals working in the architectural space and on the general public who are consuming the images? Is hyperrealism hindering the architectural rendering profession, or is it attracting attention that the field would otherwise be missing? In this thesis, I aim to investigate what hyperrealism means in the architectural rendering space, whether or not it is a welcome addition, and the impact it is having.

Architectural rendering is a field that has progressed rapidly in the last couple of decades due to the advancement of hardware and software capabilities. With the practice being so heavily tied to the hardware and software used, the constant advancements in the technology industry have seen architectural rendering make leaps and bounds in recent years. These renders are a powerful tool. They allow users to accurately visualize proposals and projects on an immense scale, both in monetary value and sheer size. The ability to create photorealistic representations of buildings of all sizes lets designers more accurately plan, get a better sense for the overall flow and feel of a project and make changes where necessary, all without having to invest any physical resources. Not too long ago, without the ability to read an architectural draft, investors, consumers and the general public only had architectural sketches to go by. These sketches were often at the mercy of the artist drawing them, meaning that they were often stylized and not a completely accurate representation of the original plans. Taking considerably longer when compared to a modern render and often times lacking in detail comparatively, the computer-rendered image has been a welcome solution to these problems. However, the relative ease with which people are able to create these digital scenes gives them the opportunity to embellish the scene in ways that would not have mattered previously. With the increased photorealism that modern technology is allowing, the ability to create a false skyline or unrealistic lighting is resulting in an increasing amount of hyper-realistic renders that previously would not have been considered a problem.

Introduction to Hyperrealism

Traditionally, hyperrealism was a movement that stemmed from paintings and was first coined by a Belgian art dealer named Isy Brachot (Conte, 2018). Brachot titled a major exhibition he hosted in his gallery “hyperrealisme”, French for hyperrealism. The exhibition featured photorealist artists from all over Europe and America and is the first case of painters referring to work as hyperreal. The exhibition contained a mixture of both paintings and sculptures that showcased hyper life-like details that had rarely been seen before.

Nowadays, photorealistic renders are becoming more commonplace. As both computer hardware and software have advanced, the technology has become more accessible to rendering firms as well as to the public. With the use of techniques such as Photogrammetry and Light Detection and Ranging (LIDAR), achieving a basic photorealistic scene now requires less effort and user input than previously required.

Photorealism and hyperrealism share many similarities, so much so that they are often confused and used in place of each other in many instances. Both methods are aiming to create or recreate a scene that becomes indistinguishable from a picture to the naked eye. Where photorealism is trying to achieve this by getting each detail exactly the same, as if a photo had been taken, hyperrealism takes this a step further and aims to achieve the photorealistic look and then add elements that would not occur organically given the context of the image. Hyperrealism is trying to achieve a look that goes above or beyond photorealism.

The range of definitions of hyperrealism depends on what context or medium the viewer is looking at. It is often defined as an artist taking their work beyond photographic quality by placing added focus on visual, social and cultural details that we perceive in everyday life (Amoakoh, 2020). This is often seen as the artist adding their own emotional opinion to a piece of work. That could be in the form of removing various blemishes to make the work more attractive or adding elements to the scene that would not naturally occur if it was to happen in real life. This also includes any other addition or occlusion to a scene that couldn't be photographed in nature.

Within Computer Generated Imagery (CGI) and architecture in particular, hyperrealism and photorealism are often used interchangeably when talking about any photo or piece of work that has been edited or enhanced to make it more visually appealing. If the edits to an image are within the realms of what is realistically achievable in real life, then the image can still be considered photorealistic. In most cases an image is not considered hyper-realistic until material that would not naturally occur in that environment is introduced into a scene. This material could be something that defies the laws of physics, such as impossible lighting scenarios. Or it could be the inclusion of a cityscape horizon to make an apartment look more attractive.

The computer science field defines realism as the ability to simulate the visual appearance of any object in such a way that its computer image is indistinguishable from its photograph (Manovich, 1995). Achieving this realistic looking scene or render has been the goal of architectural renderers since computers have been able to render 3D graphics. However, with technology always improving, the goalposts for what constitutes as a realistic render are constantly moving. Previous researchers looking into creating realistic renders had to have extensive knowledge of the mathematical equations behind things, such as calculating light bounces and determining how much light absorption is happening when an object hits various surfaces. Papers such as Fasse and Paul (1995) were released with a detailed breakdown on how to go about those processes.

Not long after the inception of computer-generated 3D images, magazines started to publish articles and research papers claiming how technology had reached the point where the computer-generated images were indistinguishable from that of a photograph. Papers such as Perlin (1985) go into great depth about various techniques used to render realistic CG textures using the technology available to them at the time. Nearly 40 years after Perlin's paper, knowledgeable users are able to create photorealistic renders using modern 3D modelling and rendering software packages without being required to know the mathematics and principles driving the software.

Historically it was thought that synthetic photographs generated by computer graphics could never be as precise in rendering visual reality as images obtained through a photographic lens (Manovich, 1995). This was a common viewpoint in the mid to late 1990's as computer generated imagery started to become mainstream in film, gaming, photography and any other industries that commonly uses CGI. However, with the introduction of ray tracing, a more sophisticated way of calculating light bounces, in 2007, rendering engines have made massive breakthroughs when it comes to replicating the way real life looks and moves. This means that one can comfortably say for the first time, that computer generated imagery has the ability to be indistinguishable from real life (**Figure 1**). One of the key fields to adopt these new lighting and rendering techniques was those involved with architectural rendering.

Figure 1: *An example of how a modern engine calculates realistic light. (https://www.unrealengine.com/en-US/unreal-engine-5)*

Chapter 2. Hyperrealism

Hyperrealism in Architecture

Architecture and architectural rendering seem like the perfect fit for hyperrealism as advancing the technology and trying to create more realistic renders seems to be the natural step in progressing the practice. But who are these hyper-realistic images being created for and to what purpose?

With the ability to produce renders to a hyper-realistic level, many professionals in the architecture industry are starting to question what these renders are worth. This leads to further discussions about the role of realistic architectural renders in general. Who benefits most from this technical and somewhat lengthy process? Are the renders an instrument to help assist architects in their design process? Do the images better help the general public understand the nature of the architectural drawings? Or are they just a visually appealing way to promote proposed buildings to clients?

Much like the postproduction pipeline for photography, hyper-realistic rendering creates a somewhat 'beautifying' effect that can transform an image based on perceived emotion or beauty standards (Moreira, 2021). It is because of this effect that the industry has started to observe some of the top architecture firms argue for the removal of high-end hyperreal renders from their pipeline. These renders are creating a sense of over sensationalised connotations within the general public and when the finished product does not match the high standards set by the original render, the architectural firms are the ones receiving the full force of the feedback. This has led some firms to exclude hyper-realistic renders from their pipeline entirely, instead choosing to favour line drawings that are more akin to more traditional forms of architecture.

Alternatively, both, smaller architecture firms and architecture enthusiasts, are generally favouring these hyper-realistic renders, because they are a way for artists to help architects (and others) to visualise an idea, not a reality. The majority of people are able to discern the difference between the two, regardless of how realistic the images are getting. For many, these rendering tools are just tools to help support a design and most people know that the result of a rendered design will not be completely representative of an idea. However, it is still clear that not everyone is able to tell the difference between reality and hyper-realistic renders, so the ability to both intentionally and unintentionally fool people still exists.

So why is there now a surge in hyperrealism and its related fields? Over the last 17 years there has been a relatively slow, but steady climb in web based searches when looking at the term hyperrealism (Trends, 2021). This may indicate the general population is not aware of the term ‘hyper-realism’, or simply that they are not interested in researching it. Alternatively, when looking at trends in academia and focussing on publications (such as books and articles) usage of the term ‘hyper-realism’ increases in the in the late 1970’s, early 1980’s, around 10 years after Isy Brachot first coined the term (Ngram Viewer, 2021). Since the 1970’s, there has been a steady increase in the number of times the term ‘hyper-realism’ has appeared in publications, with it spiking recently in 2019, which is as far as Ngram is currently tracking (**Figure 2**).

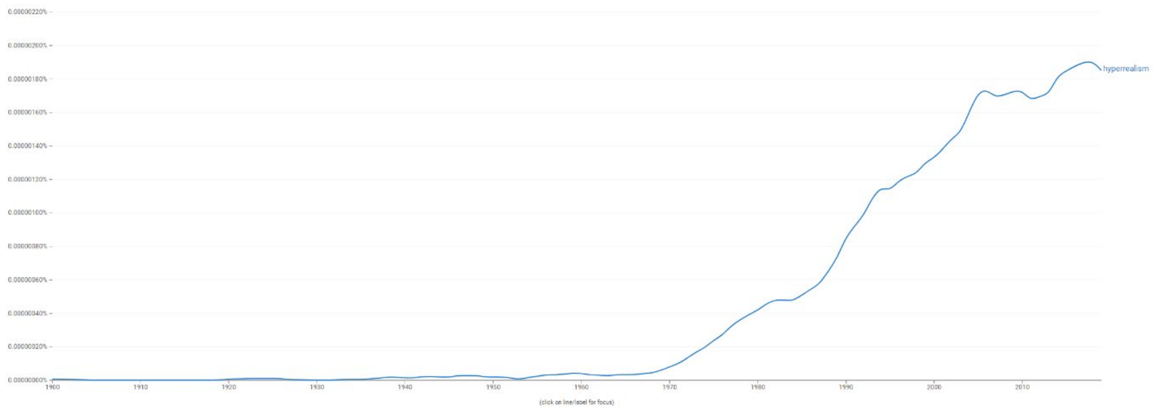

Figure 2: *Ngram Viewer Graph showing searches for Hyper-realism over time. (https://books.google.com/ngrams/graph?content=hyperrealism&year_start=1800&year_end=2019&corpus=en-2019&smoothing=3)*

It is likely that Hyperrealism is starting to be looked at with higher frequency and with more scrutiny at the research level. This strongly suggests that academics are aware of the problems associated with hyper-realistic renders and the subconscious effects that those renders can have on people.

New Vs Old

The earliest known architectural imagery can be traced back to just over 4000 years ago, with the discovery of ancient Egyptian and Mesopotamian architectural drawings (Designblendz, 2019). While rudimentary, they were considered basic architectural plans. It was not until the early 1400’s when the famous Italian architect and artist Filippo Brunelleschi first came up with the concept of the perspective drawing. A perspective drawing is defined as a technique used to create the linear illusion of depth. This concept is widely recognised as one of the fundamental ideas that marked the beginning of the Renaissance movement (Argan & Robb, 1946). Perspective drawings are one of the main ways modern architects portray their ideas to other designers and laymen alike. With over 600 years’ worth of these drawings being consumed by the public, the connotations of perspective drawing plans are recognised as just that, plans. As viewers, we understand that when looking at either a hand drawn sketch of a proposed project or a computer rendered image in a similar style, such as line art, that this image represents an idea and is subject to change. This has been ingrained into consumers after centuries of exposure to these types of images. However, this is not always the case when new forms of media are presented to us.

Unlike older forms of media, which the public have had generations of exposure to, newer forms of media such as CGI do not have these learned connotations as there has not been as much exposure to them yet. This has most definitely been the case with computer generated architectural renders and it is affecting how the public perceives them. Having prior knowledge about a certain or similar subject can influence the rate at which one can learn about a new concept (Pazzani, 1991). Lack of prior knowledge or experience with a subject can remove the ability to logically perceive and digest the new knowledge. People are able to take the prior knowledge theory (gained through experience with a subject) and apply it to the world of realistic renders. When presented with a photoreal or hyper-realistic render, most people will compare it with a photograph. Due to more than a century of exposure to photographs, it is something that people have a good understanding of and are familiar with. The photograph represents a previously existing moment in time. It is a snapshot of an event and is usually seen as fact. While photographs have always had the potential to be edited, until recently, with the inclusion of photo editing software like Photoshop, it required a good understanding and the equipment required to develop photos. This meant that they are generally perceived as factual, which can cause a few different problems when viewing architectural renders.

Because of the automatic instinct for many to perceive a render the same way one would a photograph, the average person can struggle to disassociate a render from real life if it is bordering on hyperreal or photorealistic. This can lead to high quality renders being viewed as factual. Unlike the photograph, which is generally seen as a factual representation, 3D renders can at best, only be a representation of a real scene and more often than not, they are completely fabricated and at the discretion of the artist. This can lead to a significant problem when people who are unfamiliar with the concept of lifelike 3D renders think they are being misinformed or lied to when final products do not look exactly like the render. This has been an increasing problem in the architecture and real estate industries in recent years. In response, architecture firms and those advertising through the use of renders are starting to place disclaimers on images stating that they are artistical renders and not representative of the product. As the general population gains more exposure to this form of media, this problem should occur less frequently over time.

Photography and Architecture

Photography and architecture have always gone hand in hand. In fact, the first photograph (Figure 3), taken in 1826, that we have a record of today was of someone's home (Clarke, 1997). Requiring long exposure times and minimal movement meant that buildings made for the perfect subject material for early photography. Architectural works cannot be accurately described via word of mouth or through text, without the ability to photograph buildings, centuries of work would have the potential to be lost without any form of record.

Joseph Nicéphore Niépce was a French inventor. Among other achievements, he was credited with the invention of photography and would continue to be a pioneer in this field. While the exact date of the first photograph Niépce took is uncertain, the first that is still in existence today was taken in 1826 from the view of his window. It is estimated that the exposure time required for the photo was from 8 hours up to several days.

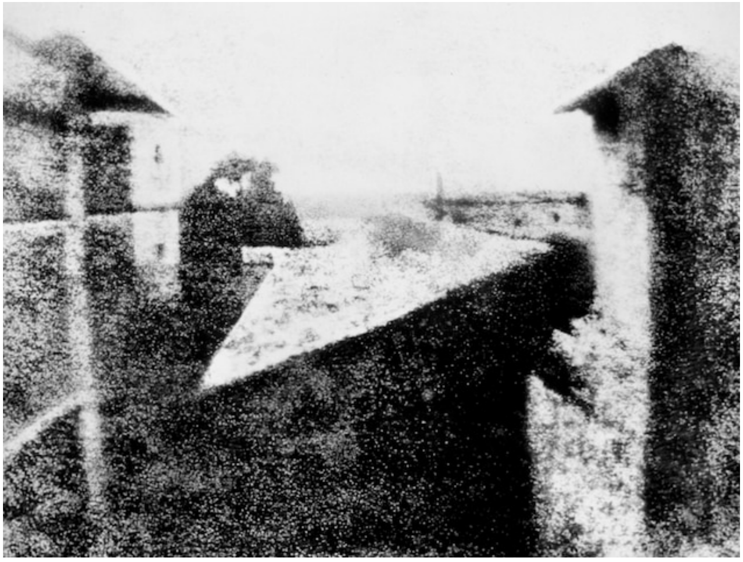

Figure 3: *View from the Window at Le Gras, Taken by Joseph Niepce in 1826 (https://www.nationalgeographic.com/photography/article/milestones-photography)*

With over a century of photos being taken and ingested by the public before somewhat realistic computer-generated images were able to start being produced, the look and feel of a photograph was already engrained into the public. As a result of this, looking at an architectural render that is fully in focus, without any form of film grain or blemish, instantly stands out to someone with experience in computer generated imagery. It is these characteristics of cameras that end up influencing the final features of architectural renders (García Crespo, 2021). Because of this, rendering software packages place a large emphasis on being able to accurately emulate the standard functions and characteristics we would find on a "physical" camera (such as aperture, depth of field, film grain, etc). It is the complete control over these functions that enables renders to be hyper-realistic as these digital images create an unrealistic, almost too perfect image (Figure 4).

Figure 4: *An example of hyperreal rendering. (https://www.archdaily.com/960849/who-are-we-making-hyper-realistic-renders-for)*

Photographic Deception

Photorealism is defined as creating an image that is indistinguishable from a photograph (Manovich, 1995) with a photo being considered a brief snapshot into reality. While it was possible to edit an analogue photo to a certain extent, it required a good understanding of the photo development process as well as the skills necessary to develop film. However, with the shift to digital, the photograph is now able to be easily manipulated with far more ease, adjusted and re-rendered (Morrison & Skjulstad, 2010). Users can now edit images in a way that extends beyond photorealism and starts to cross over into the realms of hyperrealism.

This digital manipulation is not always able to be perceived by the average person. Lange (2001) shows through extensive blind testing that 3D representations of landscapes were realistic enough that even locals to the area had trouble distinguishing between real and rendered images. While these rendered images are not inherently a problem, digitally fabricating an image in order to deceive an audience either intentionally (or by accident) is a real possibility (Mitchell, 1994). This is because the observer is relying on their prior knowledge to interpret the imagery.

"An image that claims to accurately represent a projection of a particular future in a particular setting is primed with the capacity to deceive an audience who has unwittingly suspended disbelief" (Mitchell, 1994). It is through the nature of statements like this that CGI visualisations have the ability to be so effective. Whether or not the creator is doing so intentionally, by relying on the connotations that the photograph has amassed over time, hyper-realistic visualisations are able to trick many into thinking the image represents reality. In 2013, Peter Guthrie, widely regarded as one of the leading artists in hyper-realistic architectural visualizations, was interviewed by dezeen.com about hyper-realistic renders (Ross. Bryant, 2013). He believed technology had reached the point where renders were indistinguishable from photographs, stating that "The 2013 Ikea catalogue has a surprising number of visualisations in it and most people are none the wiser." Now in 2022, with nearly a decade of both hardware and software development, the quality of computer renders has increased greatly. It is now easier for designers to convince the average person that a render represents reality even when the picture contains unrealistic elements or aspects such as temporal conflation.

Temporal Conflation

Through the nature of computer-generated images, artists are able to add physical elements into a scene at any point in their existence. When each element of a render is added at their most attractive moment in time, this leads to a concept called Temporal Conflation (Kullmann, 2014). For example, by using trees that are 20 years old, wood with a 5-year-old patina and uncorroded metals, coupled with lighting that is generally unachievable naturally, the artist is inadvertently creating a hyper-realistic scene of a building that misrepresents what is feasibly possible. This only further creates unrealistic expectations for potential buyers and thus can lead to disappointment when the final product cannot meet the standards set by the renders.

Rendering an image or taking a photograph does either one of two things when it comes to time. First, it can take a snapshot of a time period, rendering the subject "prisoner" of a particular moment in time. This is particularly true when it comes to images that contain subject matter we can easily date, such as cars, peoples' clothing or technology present in the scene (Borden, 2007). This can be a useful technique, especially when the photographer is trying to portray a certain message through their imagery. Alternatively, it can do the opposite and erase time all together when a scene does not contain content that can easily date an image, such as clothing or vehicles. The latter is what is generally seen with more traditional architectural photography. Photographers are often trying to capture a building in a timeless manner, aiming to document a structure without external influence (**Figure 5**). "Take those pictures of buildings caught perfectly before people, dirt, rain and history move in: since the beginning of the twentieth century it is these pictures which have framed a history of architecture in both its production and reproduction – a history, in which architecture is seen to be a stable power, existing over the dynamic forces of time" (Borden & Rendell, 2000).

Figure 5: *Architecture from the 1950's. Focusing on natural elements to create a timeless design. (https://blog.buildllc.com/2011/05/what-does-it-mean-to-be-timeless/)*

While the inclusion of time-relevant objects is a tactic that is often used to emphasize a particular historical period in architectural imagery, the overall aesthetic of the image can often have the same effect, either consciously or subconsciously. The colour of an image, the overall quality and composition all affect how a viewer can depict the age of a photo. This is particularly true for black and white images, with many people making the assumption that a photo is considerably older solely because the image is monochromatic.

However, the opposite can be true as well. Images released by an anonymous photographer who used the Dufay colour process (a process in which a photograph is viewed through a series of mosaics containing small colour filters to artificially give a monochromatic image colour) showed images of bombed houses in London during World War Two. Viewers said that the photos seem “unreal” and “inauthentic” since they were polychromatic. Because people usually associate photos of that era with black and white, seeing the photos colourised automatically leads them to think they are not factual (Seaborne, 1995). We also see similar results when people are viewing photos with a filmic grain or distortion applied to them. When applied well with relevant subject matter, such distortions can subconsciously trick the user into thinking they are viewing older images, which people assume to be authentic, due to older photography being so much harder to edit.

While temporal conflation is a principle that occurs quite often, it is reasonable to assume that many designers are creating images while being completely unaware of this influence. For most, an architectural render is created with the purpose of representing a product in the most aesthetic way they can achieve to increase its saleability. Whether knowingly or unknowingly, they will combine materials and objects of varying ages and origins to achieve the look they desire causing temporal distortion. However, the combinations of these elements do not always have the desired effects on the viewers. When done poorly they have the potential for the render to start bordering on the uncanny.

Uncanny Valley

In 1970, a robotics teacher named Masahiro Mori, who was working at the Tokyo institute of technology wrote a paper on how he thought people would respond to a humanlike robot. Mori hypothesized that a person’s response to humanlike robots would very quickly shift from empathy to revulsion the closer the robot resembled a lifelike appearance without actually reaching that goal (Mori, 1970) This phenomenon is known as the Uncanny Valley (**Figure 6**).

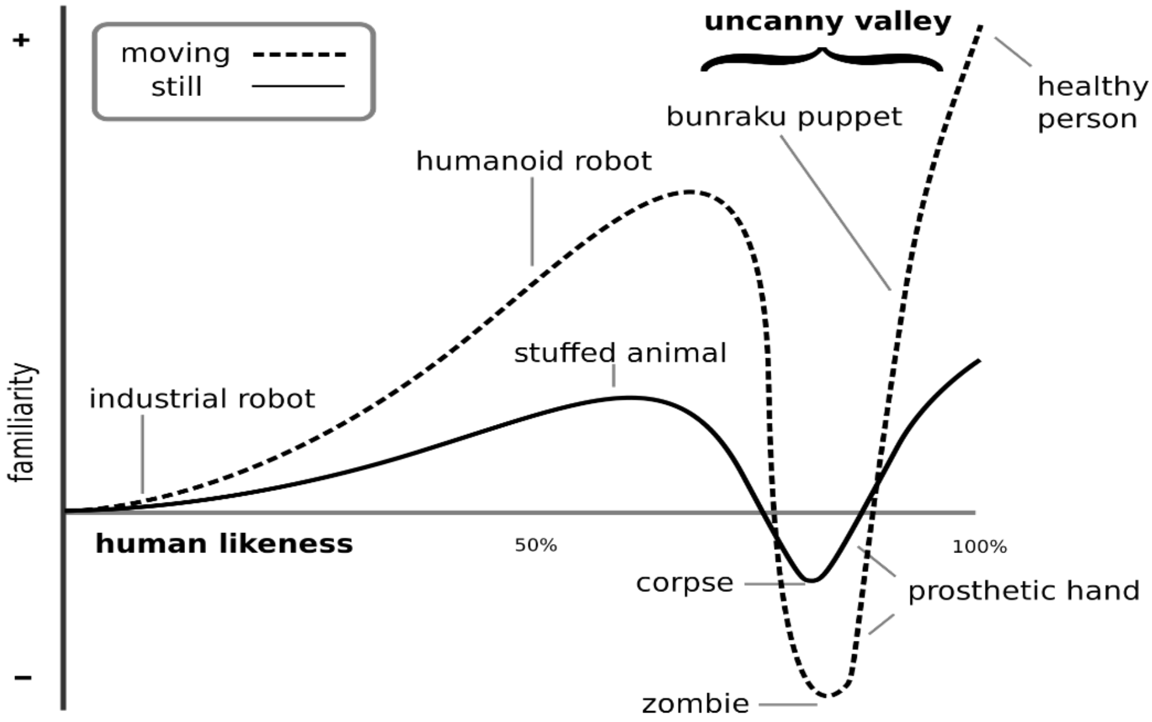

Figure 6: *Mori's hypothesis on the Uncanny Valley (Mori, MacDorman, & Kageki, 2012).*

Mori's original writeup on the topic was published in a small journal in 1970 and did not garner much attention at first. However, in more recent years, the concept of the Uncanny Valley has gained interest with an increase in publications and websites talking about the phenomenon as more CGI is being released (Mori et al., 2012). As technologies have evolved, scientists working in fields such as human-robot interaction and computer graphics are getting closer to reproducing humanlike facial features and movements. The notion of the Uncanny Valley has become as prevalent as Masahiro Mari first envisioned. However, not all agree with Mori's theories around the Uncanny Valley. Teams such as Hanson Robotics Inc based in Arlington, Texas, argue that industries should be aiming to create as lifelike as possible recreations in the various scientific fields, as it helps with human perception and cognition (Hanson et al., 2005). The team is even going as far as conducting multiple web surveys with the intention of proving that the general population is not concerned with the visual aesthetic when approaching, but not quite achieving a lifelike representation. While Hanson's study does look into whether or not the public find the varying levels of realism 'acceptable', it does not investigate or compare how well a product will perform at market when comparing the varying levels of realism to each other (Hanson et al., 2005).

Mori's theories around the Uncanny Valley in robotics and other sciences is still applicable when looking at technology driven fields such as architectural visualisations, albeit not quite to the extent that he first envisioned. Architectural visualisation photos that are striving for a photorealistic or hyper-realistic look but manage to just fall short generally do not instill a sense or revulsion or disgust. This is generally because without a human subject, people do not find the images to be uncanny, instead they tend to just break a viewers' suspension of disbelief. Where architectural renders do tend to fall into the realm of the Uncanny Valley, is with the inclusion of people.

Like many of the other problems that stem from computer generated imagery, the Uncanny Valley effect is driven purely by the way humans perceive what they are looking at. Humans are very habituated to the way people move and how they physically look, to the point where if a computer generated (CG) human looks and moves slightly unrealistically, people will instantly notice. When it comes to creating both photorealistic and hyper-realistic imagery, trying to replicate human intricacies in detail takes large teams of people working hundreds, if not thousands of hours and they are often still not capable of accurately recreating human movement. A rather famous example of this is Steven Spielberg's "The Adventures of Tintin" which was released in 2011 by Peter Jackson and the team at Weta Digital (WetaFX, 2022). As Kyle Buchanan, a reporter for the New York times said "There's something off about this character, and it is clear why he has been minimized in the trailer for his own movie. Aside from the swoop in the front of his hair that lends him some cartoonish verve, Tintin looks simultaneously too human and not human at all. His face weirdly fetal, his eyes glassy and vacant instead of bursting with animated life" (Buchanan, 2011). (**Figure 7**).

Figure 7: *Steven Spielberg's The Adventures of Tintin (https://www.salon.com/2011/12/20/the_adventures_of_tintin_spielbergs_weird_action_cartoon/)*

When dealing with animated characters, most films will choose to add a layer of caricature or stylised animation to their characters. This clearly indicates to the audience that these are not actual portrayals of humans and immediately removes any notion of the Uncanny Valley, such as Captain Haddock's comically large nose (**Figure 7**). When creating Tintin for the 2011 film, Peter Jackson and the team at Weta Digital chose to employ the latest techniques in facial and body capture. They aimed to make Tintin, a comic book character, resemble a lifelike person as much as possible. However, because every other character was slightly caricaturized and as a result of the facial motion capture Tintins eyes looked flat and lifeless, often making it look like Tintin was staring blindly into the distance. Tintin stood out as this hyperreal figure that left the audiences feeling uneasy (Kunz, 2015).

The Uncanny Valley effect shows itself differently in the architectural rendering space. Some 3D artists use a stylised look to portray their work, but the majority are aiming to create scenes and images that are as photorealistic as they are capable of producing, to give the viewer the most immersive experience possible. "In conjunction with a render that is 'too perfect' entering the Uncanny Valley, the presence of too many fine details being in focus in a render can also leave the viewer of a render with a similar Uncanny Valley effect as they get overwhelmed with the details in the image" (Maganga, 2021). This effect is further amplified when trying to mimic complex architectural scenes that contain many different surfaces with varying textures and reflective values. This effect can be lessened when the designer is actively creating scenes that are clearly fantastical and would not be found in a real world examples, such as Alexis Christodoulou's renders in which he takes inspiration from various paintings and abstract modernism (Santibañez, 2019). With the inclusion of elements such as dark volcanic rocks or perfect geometric shapes places in his scenes, it makes it much easier for the viewer to determine that they are computer generated renders.

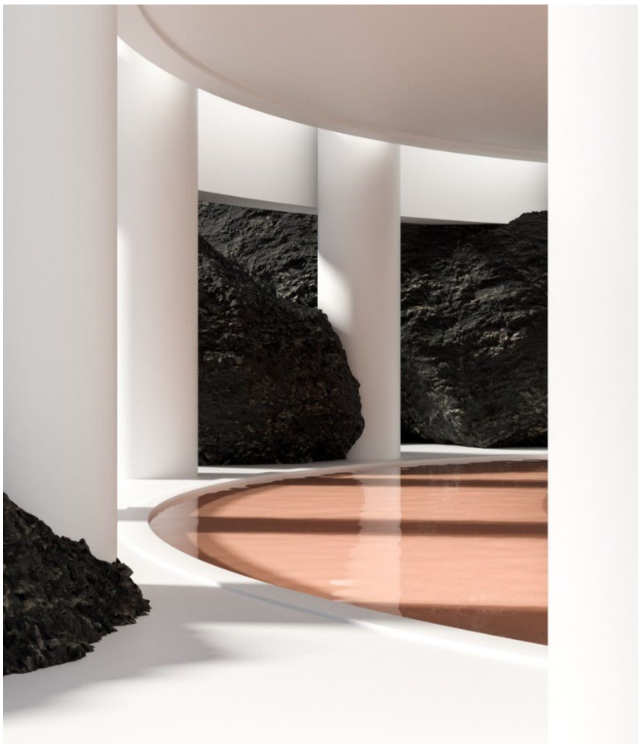

Figure 8: *Alexi Christodoulou Hyperreal Renders (Santibañez, 2019).*

Renders that are striving for photorealism but come up short are often considered to fall into the Uncanny Valley. However, this is often due to the viewer not being able to determine what is wrong with the image. Where architectural renders really start to give viewers that uneasy feeling is through the inclusion of life-like figures. Older forms of architectural rendering such as montage, have been using images of people for nearly 100 years without a problem. But due to the nature of that technique not aiming to create a photorealistic scene, there has not been problems with it looking unusual because it does not fall into the uncanny valley. However, newer software packages are now giving users the ability to add 3D models of people into virtual environments with the click of a few buttons. Due to the nature of trying to recreate a human in 3D, these models generally come out at a much lower fidelity compared to the rest of the render. This results in lifeless, almost plastic looking human figures sticking out in these hyper-realistic architectural scenes, and this fits perfectly into Mori's definition of the Uncanny Valley (**Figure 6**). As Henry Goss, who is considered to be one of the best modern architectural renderers in today's society said "There are two main ways of adding people to renders, either by rendering a 3D modelled person or by montage in post-production. The human brain is highly tuned to pick up the subtle nuances of human appearance and movement and therefore it is very difficult to achieve a convincing result" (Gavrila, 1999).

Figure 9: *An example of Uncanny people in a hyper-realistic environment. (https://student.aut.ac.nz/social-responsibility)*

Having an awareness of what physical elements in a scene can cause the Uncanny Valley effect and how they can be avoided is an essential skill when creating an immersive render. Being able to determine when to fill a scene with 3D models or when to take a more minimalistic approach can greatly aid the designer in making a convincing environment. "As with a lot of architectural photography, sometimes the pure nature of the architecture itself is better represented laid bare" (Ross Bryant, 2013). Different methods of model creation and acquisition can result in vastly different results when creating a 3D environment. Knowing the strengths and weaknesses of the various methods and when to apply them can be a great addition for any designer creating realistic scenes. The following section introduces the two most common acquisition methods. Photogrammetry and LIDAR.

Model Acquisition Methods

Photogrammetry

Photogrammetry is a process in which a digital model is composed based off data derived from a series of photographs. Photogrammetry as a concept, dates back as far as 1850 where it was being used to help map landscapes using perspective-based photos. World War I greatly helped to advance the technology as there were tactical advantages to being able to map terrain based on ground and aerial photography (Gosh, 1981). However, it was not until the late 1990's that the first examples of Digital Photogrammetry started surfacing (**Figure 10**).

Figure 10: *A house constructed with photogrammetry (https://bitfab.io/blog/photogrammetry/)*

Debevec, Taylor, and Malik (1996), a team working out of the University of Berkeley CA, were among the first to create a purely image-based photogrammetry algorithm. Previous iterations of photogrammetric renderers used a hybrid approach and geometry-based methods. Comparatively, these two methods were considered to be far too labour intensive, as the user had to manually register the images and create tie points between them. Not only did this take a considerable amount of time, but it is also far less accurate than computer generated tie points, often resulting in a final model that has taken longer to generate and is less accurate. The image-based method where relatively complex algorithms can match a set of photos together without user input is still widely used to this day and is capable of producing high quality 3D models.

As the mathematics behind the software has become more robust and the software has become more user accessible, more people at both, the consumer and business level, started using photogrammetry. While previously a large amount of photogrammetric projects were led by research teams, more disciplines outside of academia started to adopt the technique to help with many aspects of their work, including digital preservation and conservation, historical documentation, digital monitoring of assets, virtual reality and computer generated graphics applications, among others (Remondino, 2011).

It is through this new photogrammetric process that users are able to recreate scenes ranging from small objects to entire buildings, both interior and exterior. (Shashi & Jain, 2007) show that a single consumer grade Digital Single-lens Reflex camera (DSLR) that has gone through a basic calibration process can be used in the photogrammetric processes and result in a 3D model that is as accurate as other competing computer-aided design (CAD) methods. It is both faster and requires less equipment than previous photogrammetry methods (Shashi & Jain, 2007). However, as the scale of a project increases, more expensive equipment is required.

Where photogrammetry falls short, and other methods become more favourable, is when dealing with very intricate surface geometry or reflective materials. One of the larger downfalls of the photogrammetry process is that when dealing with an untextured or reflective model that has a lot of intricate details and height variation, photogrammetry can create physical artifacts in the surface geometry, in the form of valleys and peaks (Dixit, Kennedy, Piemontesi, Kennedy, & Krebs, 2019). As a result, the amount of manual clean-up required greatly increases the length of the workflow.

Even though photogrammetry can lapse in quality when it comes to very intricate and reflective objects, it compensates for these downfalls by being able to create highly detailed texture maps. 3D photogrammetry generates its objects based on pixel information. This means that texture maps are created pixel by pixel based on the 2D images that were used to structure the 3D model (Dostal & Yamafune, 2018). Since digital photos are used as the source data, the quality of the photos themselves heavily impacts the final output. The quality of the camera, how skilled the photographer is, and if applicable, the environmental constraints like shadowing and lighting all impact the final result. Previously, photogrammetric software packages also required manual input of the camera specifications, such as focal length, aperture, and exposure. If these did not match the metadata being fed into the software or if there was variance in the photos received then the final output would suffer greatly (Green, Matthews, & Turanli, 2002). However, with advancements in the software since the early 2010's, the need for pre-calibration of cameras is no longer required, eliminating a very time consuming and error prone part of the pipeline.

With the more recent advancements in the software and the ability to use consumer grade cameras, there has been a sudden surge in popularity of photogrammetric objects. This is contributing to hyperrealism in architectural renders. Previously, organic shapes such as furniture and assorted items commonly found around a room had to be hand modelled, individually. This generally resulted in 3D models that were not completely true to their reference, whether the scale was slightly off, or the subtle curves of the object were not accurate. The lack of these small natural irregularities greatly reduced the level of immersion that an average viewer would experience, which helps them to distinguish a render from a photograph. With the increased accessibility of photogrammetry in today's 3D pipeline, many of those small irregularities that would previously have occurred are removed during the process. This is resulting in much higher levels of immersion due to the nature of the photogrammetric process.

LIDAR

LIDAR, which stands for Light Detection and Ranging, is a tool that was created under the direction of the Hughes Aircraft Company in the early 1960's with the original intended use being satellite tracking. By emitting laser impulses and calculating the time it takes for the laser to return, the system is able to determine how far away an object is from the laser's origin. All LIDAR systems are still based on the same technology. By shooting out millions of rays and calculating the distance of each reflection, LIDAR is able detect the varying surfaces of an object and compile this data to build a very accurate 3D model with accurate depth parameters. Laser scanned data often tends to produce more consistently accurate dimensional data when compared to other digital documentation models such as photogrammetry. However, laser scans do fail to produce the same high quality and high resolution textures that correctly convey colour information in comparison to photogrammetry (Dostal & Yamafune, 2018). Another disadvantage of the LIDAR scanner itself is that it needs to be fixed to a solid structure and cannot be moved during the scanning process without serious data degradation. Often, multiple scans have to be performed to make sure that adequate coverage of an object is achieved and there are no occluded areas. The multiple scans need to be manually aligned and compiled to create one solid 3D object. This is a lengthy and complicated process that does not always have the desired results (Sapirstein, 2018). Without the use of polarising filters to manage reflections and suppress glare, it is difficult to capture both reflective and transparent surfaces with LIDAR., These surfaces often need to be covered with a matte material, such a patterned fabric or newspaper, and then manually recreated by the user at a later stage. Newer lightweight, handheld LIDAR scanners are now becoming more available and are capable of producing quick and easy results without some of the downsides of fixed-in-place scanners (Bauwens, Bartholomeus, Calders, & Lejeune, 2016). However, they still lack the accurate depth information that a fixed LIDAR scanner can produce.

Due to both the positives and negatives of both photogrammetry and LIDAR, some people have adopted workflows that contain information from both methods (Markovic, 2018). Combining the accurate geometry created by LIDAR scanning with the accurate colour data that photogrammetry provides, a user is able to create a well-rounded 3D model that consists of accurate topology as well as colour data. The photogrammetric process is also excellent at gathering coverage of the areas that LIDAR would have trouble reaching, such as coverage behind larger objects and occluded areas. This process does however require a good knowledge of how the LIDAR equipment works and what will and will not be captured in the process. While using both LIDAR and photogrammetry together can result in a higher fidelity model, the post processing workflow is much more cumbersome than working with either one individually. There is a greater potential for the two data sets to be corrupted when combining them and it requires a highly skilled user to create a successful final product. When the two are combined correctly, there is the potential to create a model that has both higher fidelity surface geometry and colour information with less information missing. An example is presented in **Figure 11** below.

Figure 11: *A combination of both photogrammetry and LIDAR for a high-fidelity model (https://danielcahill.artstation.com/projects/yOvBn)*

How to make a hyper-realistic image.

The last decade or two have seen a surge in the development of 3D graphics and the computer rendered image and as a result creating hyper-realistic architectural renders is now easier than ever. Having the knowledge of what elements are used to create a hyper-realistic render can help people to discern between photoreal and hyperreal.

There are common mistakes that people often make when it comes to trying to create realistic scenes that can quickly break that sense of immersion a viewer has, greatly detracting from render. The first of which is that the potential for hyperrealism is “limited by a rendering artist’s technical capabilities and their eye for a well-crafted composition” (Keskeys, 2022)

As previously addressed, when discussing old vs new technology, there is an emphasis on the importance of the camera and the photograph on how people perceive images. Trying to emulate the qualities that a camera lens introduces into digital renders is one of the most impactful factors when it comes to differentiating between a photograph and a rendered image. The most obvious of these factors is replicating focus. Without any adjustments to the camera within the modelling or rendering software, every object in view will be in focus. This is a common mistake made in architectural renders. Without a foreground, a midground and a background, the rendered image conveys no sense of depth. This use of depth of field is a technique used in almost all forms of photography. When used correctly it results in a more realistic reproduction of the subject (Weber, 1980). When viewing a photo, the human eye is naturally drawn to the in-focus subject. When an image encapsulates a scene that should have varying levels of focus, but all is in focus, our brains instantly recognise that something is not right, it can even start to border on Uncanny Valley territory (Mori et al., 2012).

While focus and camera aperture is often seen as a make-or-break point for hyper-realistic 3D renders, there are other factors that also play a role in turning a high-quality render into something photorealistic or hyper-realistic. By design, 3D development platforms and the majority of other math-based software tools create rather symmetric and clean designs. However, at a fundamental level, asymmetry and imperfections are an integral part of nature (Baianu, 2012). This leads to another common mistake when trying to create hyper-realistic renders: making the scene too clean and too perfect. Less experienced designers are often hesitant to create imperfect scenes, but humans live messy lives which results in messy photographs. Having all the chairs around the dinner table perfectly aligned with a placemat neatly positioned in front of them is not representative of how most people live. Adding imperfections to a scene makes it more relatable to a viewer. In terms of architectural rendering, it makes the scene feel more lived in and better helps a viewer to picture themselves occupying the space. Fine details, such as crumpling up the rug on the ground or making the pillows on the couch unevenly spaced can improve the how realistic a render appears.

One of the harder physical qualities that 3D renderers have had trouble recreating since their inception is accurate lighting. A major problem in realistic imaging is matching tone-reproduction (Ferwerda, 2003). Lighting from both natural and man-made sources create vast ranges of lighting coverage that is unique to any physical space, resulting in great variances in how the human eye perceives colour, depth and shadowing. A computer-based lighting algorithm requires many different mathematical lighting equations and intense processing power to replicate these complexities. Even if the 3D renderer can calculate these factors correctly, it still requires the user to have a solid understanding of digital light placement, to understand appropriate lighting levels and colours, and what lights to use for various scenarios. Without a deep understanding of lighting techniques and the software to correctly reproduce them, 3D scenes are often either under or over exposed, with incorrect shadow depth and lighting. This creates an image that is obviously rendered and far less of an effective representation because it is not realistic.

Similar to how many designers make the mistake of creating scenes that are too perfect, many 3D modellers create geometry that is too square. By default, 3D modelling software packages create objects that have perfect 90-degree edges. This results in hard surfaces like tables or countertops having pristinely hard and sharp edges, whereas in reality, even objects that meet at a right angle often have soft, rounded edges. Adding a gentle bevel to objects can have the simple effect of reducing the boxy looking objects that modelling software creates by default, which can greatly enhance the effect of realism designers are trying to achieve. The same goes for adding smaller details, such as surface wear and scratches to surfaces made of materials like metal and wood. It is very rare to find these materials in the real word with no level of wear to them.

Modelling objects to the correct scale is another factor that is often overlooked and can result in skewed scenes. From personal experience, many downloadable assets can be slightly off on one or all of the axis. Building a scene based around these objects can result in items that are proportionally incorrect. This is especially problematic when working with scenes that contain larger, boxier models, such as when working with the exterior of buildings. These distorted scenes can often give the impression that the viewer is looking through some sort of warped lens.

Following the theme of camera angles: where a user positions the virtual camera in a scene plays a major role when it comes to emulating photography. A user working in a 3D modelling package is not bound by the laws of gravity a normal photographer would experience. They are able to move all around the scene, placing the virtual camera at any point in the space to capture the view that they desire. This has both positive and negative effects when trying to create photorealistic and hyper-realistic pieces of work. An example of a positive effect is being able to use an unusual perspective which is attention grabbing for the viewer (Weber, 1980). However, images that are too unusual may be rejected by the viewer if they are not understandable (Weber, 1980). The same photography concepts, including placement, angle and framing, that (Weber, 1980) talks about in his book are even more applicable to virtual cameras in 3D space. Having full control of camera placement, aperture and focal length allows the user to get very creative with how they frame their shot. However, when camera placement starts to stray too far away from the natural eye level of normal photography, as seen in (**Figure 12**), is generally taken from, photos start to look odd due to seeing objects and creating eyelines that people are not used to seeing. When shown perspectives that are out of the ordinary, they are often then associated with computer generated imagery. Keeping the virtual camera at or around eye level allows for a much more realistic perspective, further enhancing the photorealism aspects.

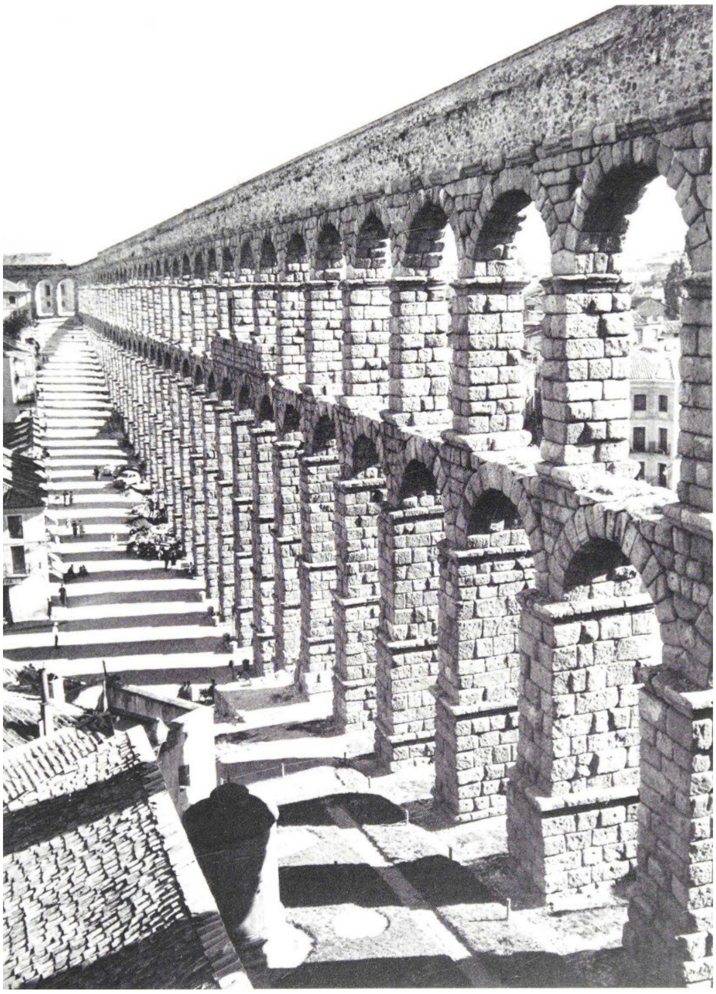

Figure 12: *An example of a photo taken from a more aerial perspective creating different eye lines. (Weber, 1980)*

The Technology behind the render

What technology is available and how that technology is being developed will always be one of the main driving factors for the advancement of 3D rendering. Since the mid 1990's, people have been quoted saying that computer graphics have reached their peak and that they are indistinguishable from real life. However, it is easy to look back retrospectively and see that it was not always the case. Renders are now at a point where a stranger can be shown an image created by a talented 3D artist and a photograph of the same space, and they will have trouble discerning between the two. This can largely be attributed to the capabilities afforded to us from modern technology.

The three main aspects which contribute to the overall quality and fidelity of a computer-generated render are Texturing, Lighting and 3D Modelling. These three factors have various technical aspects behind them and have been progressed to be able to create hyper-realistic renders.

Texturing

One of the more important properties of an object that directly contribute to its appearance and how a viewer determines what the object is made from, is the material. However, the ability to edit the materials of an object so that said object can accurately represent their intrinsic physical properties is not an easy task (Liu, Ceylan, Yumer, Yang, & Lien, 2017). With modern software packages, the user can dissect what elements are used to make up these physical properties, break them down into various categories and have full control over them. This is achieved using Physically Based Rendering (PBR). Physically Based Rendering is a concept that has been in development since the 1980's by a team of researchers trying to establish a solid theoretical base for rendering that is based on physical principles (Greenberg et al., 1997). One of the first software packages to adopt the PBR method is now the industry standard Substance Design and Substance Painter (Adobe, 2022b), originally developed by Algorithmic and later absorbed by Adobe. As PBR started to grain traction, the book "Physically Based Rendering" was released in 2014 (Pharr, Jakob, & Humphreys, 2016) by Matt Pharr and his team, which more widely popularized the term.

Physically Based Rendering is more of a concept than a strict, defined set of rules (J. Wilson, 2022). Unlike many of the previous or more 'traditional' rendering methods that sought to separate surfaces between non-reflective and reflective, the PBR model recognizes that, to accurately recreate real world materials, every material will need to have some reflective properties. Even surfaces which are considered "flat" objects, such as carpet or concrete, will reflect at least a small degree of light. To accurately do this, the PBR workflow uses a series of texture maps that it applies to every material. Previous material systems would only apply a 'reflection map' to objects that are clearly reflective, such as metals and water. Physically Based Rendering assigns every material, regardless of the perceived level of reflection to the human eye, some amount of reflection value. One would be very hard pressed to find a material using PBR that has a reflection value of 0. The PBR method then takes these reflective values and uses them to assist Global Illumination (GI), an algorithm that not only calculates light directly from the light source but also calculates light bouncing off nearby objects. This typically results in softer shadows, glossy reflections, indirect illumination and colour bleeding of coloured surfaces onto other nearby objects due to the diffuse reflection created (Lafortune, 1996). These features help to achieve the physical properties one would see in the real world. A single, bright, indirect lighting source can comfortably illuminate an entire room, similar to how in a real-world setting, light coming in through a single window would achieve a similar result.

Older, more rudimentary methods of rendering, such as the Phong method (Blinn, 1977) only calculate what is called Direct Illumination. This is the process where shadows and reflections are calculated from a single source of light without the aid of light bounces from non-light emitting objects. This usually results in very hard shadows and rather unrealistic lighting. This means that, often times, a single light, regardless of how bright, would not be enough to light a basic room. There will be many dark spots as light is not accurately bouncing around the room and lighting the far side of objects and corners. On the positive side, Direct Illumination is much easier to calculate due to the rather simplistic models it uses, resulting in much faster renders. These faster render times mean Direct Illumination is still used in most forms of rendering that require the image to be processed in real time, such as video games. Around 2021, hardware has started becoming powerful enough to calculate Global Illumination fast enough for real time rendering. Modern software such as AAA games and 3D editing tools have options to enable ray traced graphics and Global Illumination, given a sufficiently powerful PC.

The PBR workflow generally consists of up to four different texture maps combined to make the material. The first of these maps is the Albedo map, also known as a Colour map. Albedo is the most common texture map used and defines the characteristic colour aspects of an object. Traditionally, to compensate for the lack of GI in the Phong shading model, 3D artists would store or “bake” lighting information such as ambient occlusion and shadows into the albedo map. However, these aspects are now taken care of by the graphics engine and a manual baking process would often result in a worse result. Allowing the graphics engine to calculate these values means that the albedo does not have to be designed for a singular lighting configuration, which was one of the downfalls of older rendering methods. Instead, the PBR method allows for materials to work in any scene or lighting scenario that it is placed in.

Next is the Normal map. This map is used to create the illusion of depth and fine details using vector-based data encoded in the red, green and blue values of the map (Lin, Lan, Wang, & Chen, 2014). Normal maps are particularly effective in models that have lower geometry and need to fabricate the extra geometry to give the illusion of a higher detail model. The calculations driving the normal map have not changed since the more traditional ways of texturing, however, the improvements to rendering engines means that more detail, such as better depth calculations and controlling how the light interacts with an object, can be extracted from a Normal map as well as the Normal map being more effective than previous iterations.

The third map used in the PBR workflow is the Roughness map. Often referred to as the second most important texture and artist will work with after the Normal map. The Roughness map is used to control how shiny an object appears. When an object is perfectly smooth, with a mirror like appearance, it is given roughness value of 0. When an object is dull or matte, such as paper, it is given a high value close to 1. This 0-1 value controls how reflected light rays bounce off the surface of the object. Smooth objects bounce light off with little to no interference meaning that one can clearly see reflections in the surface of the object. In contrast, rough objects disperse the light that bounces off them, resulting in a blurrier reflection. An object that is completely rough will disperse the light in all directions, to the point where the reflection is so blurry that no discernible structure can be made out.

The final map is the Specular map. This is generally seen as the most difficult for artists to understand, as it is solely based on mathematical physics concepts. The Specular map is used to control the specular reflectance, which is the light that bounces off the shiny part of a surface and works in conjunction with the Roughness map. The specular value is used as an input into the Fresnel equation (Lvovsky, 2013). This equation then calculates how much light is to be reflected. The specular reflectance of a surface is dependent on its Index of Refraction value (IOR). Objects can usually be broken down into three main groups when talking about specular reflectance: metals, non-metal objects (such as water, stone, skin), and when necessary, reflective objects that also let light through them (such as gemstones and stained glass). As previously stated, working with specular maps can be considered the most difficult to do correctly, but they are essential in achieving photorealistic reflections on surfaces in 3D renders.

Through the combination of these four texture maps, nearly every surface can be replicated through a modern rendering engine. If required, extra details can be added to a surface through the use of texture maps other than main ones listed above for example, Displacement maps, Ambient Occlusion maps, and Height maps. However, regardless of what environment or what lighting setup a user places a PBR material in, if it has been created to match the proper PBR guidelines, it will greatly enhance the look of objects.

While many designers source their 3D models and textures from various online libraries such as Quixel Megascans (Quixel, 2022) which already follow a PBR workflow, having the knowledge to be able to accurately make adjustments to the textures so they still meet real world values is a very valuable skill. Being able to this means that users are able to create virtual scenes that stand out from the rest, even when using publicly available assets.

With Physically Based rendered models it is easier to achieve photorealism with fewer but more intuitive parameters. It also allows the user to achieve consistent looks under different lighting conditions, making the materials created with PBR in mind far more applicable to multiple scenes, saving a great amount of time and effort (Seymour, 2013).

Lighting

One could argue that lighting plays the most important role in creating lifelike renders. There are many functions that lighting has an impact on in virtual environments. It helps to direct the viewers eye, it has the ability to evoke certain emotions or to set the atmosphere for a scene (Seif El-Nasr, 2005). It also has more direct outputs, such as giving environments a sense of depth and accurately displaying existing texture maps and models. An incorrect lighting setup can easily break a viewer's immersion, lead their eye to the wrong part of an image or misrepresent the setting for the image. Regardless of how well a 3D object has been modelled, without the correct lighting setup, it has the potential to come across as flat and unconvincing. Conversely, good lighting has the potential to enhance even poorly created models by placing the emphasis on the lighting and shows that are created (Eye, 2021).

Lighting in 3D environments follows similar techniques to those of photography or cinematography. Good lighting techniques that apply to those fields also work with virtual lighting, such as three-point lighting and the various techniques for creating different shadowing effects. One of the features often forgotten that greatly enhances a virtual environment is what is happening off frame. Hours are often spent curating the content that is in the frame, with little effort put into what is happening in the immediate vicinity. Unless one is viewing photos and media that have been staged in a studio setting, with full control over lighting and props, often many photos have reflected light or shadows from objects outside of the frame "spilling" into scenes. While this may appear to detract from a photo, it is often what gives a photo a feeling of being alive and not staged. The same goes for a rendered environment. Depending on the rendering engine being used, when light hits an object off frame, light bounces are still calculated based off the roughness and colour of that object. This results in different coloured light bouncing around the scene, giving it a feeling of being alive or real, as opposed to just relying on a directional light to emulate the sun or another light source. The same applies to shadows. Off-screen shadows add a huge sense of depth to a scene and aid greatly when trying to achieve a certain mood through a render.

Humans are very sensitive to the effects of lighting. It has the ability to impact our mood, circadian rhythm and our daily behaviour (Chen, Cui, & Hao, 2019). As a result of this, we are very sensitive to lighting variations and can detect when the lighting is odd in a virtual environment. This is generally most noticeable in a rendered scene when looking at shadows.

When looking at a real-life scene, no two shadows are the same. There are variances in how sharp the edges are, how dark the shadow is, and something called contact shadows, which are often found where two solid objects come into contact with each other. These are especially prevalent when looking at objects sitting on top of carpet or other surfaces that are not perfectly flat. All of these factors only get further amplified when adding multiple light sources to a scene. To be aware of how shadows react in different scenarios and being able to replicate that inside of the various 3D programs gives the user a much greater chance of replicating photorealistic lighting.

Modelling

Of the three different 3D rendering aspects presented here, modelling can be considered the most forgiving. Unlike texturing and lighting, there is a sweet spot where models that are slightly subpar have the capacity to be carried by the other two elements. As long as the object was modelled to the correct proportions, a large amount of the fidelity can be made up with good texturing and lighting.

Digital models are present in multiple forms of media nowadays. The sheer amount of them supports the illusion that 3D modelling is simple. However, the generation of a precise and photoreal 3D model of organic and complex objects still requires a considerable amount of effort and skill (Remondino & El-Hakim, 2006). By design, 3D modelling packages are very heavily mathematically based. All forms of 3D models can be broken down into raw numbers. As a result of this, without alteration, models will end up being created in the most mathematically efficient way possible. This is most noticeable when creating any object with hard edges, particularly those with right angles. When looking at surfaces in the real world that come to a right angle, most edges will have a small amount of rounding, or a bevel on them. Even when looking at objects such as machined wood, the intersection of the two faces will have a small bevel to some degree. Without artist intervention inside 3D modelling packages, the intersection of two faces will come to a razor-sharp edge. This leads many who are unfamiliar with creating life-like 3D scenes to have very boxy or square, artificially made 3D models. This is especially prevalent in the architectural industry as the modelling required for them contains a large amount of right angle and hard-edged surfaces.

While there are other aspects to the creation of digital asset creation, they generally stem from these three main principles, texturing, lighting and modelling. Having a good understanding of these allows for far greater control and customization when working in 3D software packages. When a designer can effectively combine these three elements in the creation of 3D objects and scenes, the potential for creating photoreal and hyperreal rendered images greatly increases.

Chapter 3. Hyperrealism and Architecture

The Rendered Image

The use of the rendered image to help achieve a sense of hyperrealism is not something new to computer generated images. Prior to the development of the technologies we have today that are used to facilitate digital realism and hyperrealism, the limitations of analogue and earlier digital techniques required the use of abstract and hybridized approaches (Kullmann, 2014). Using techniques such as traditional sketching, picture collage or montage and wireframe perspective images, architects and artists were still able to create both realistic and hyper-realistic images, albeit this often required significantly more time and skill to do while producing less realistic images than are able to be created today. While these techniques used to be the major drivers when it came to creating hyper-realistic images, the use of the computer-rendered image has meant that less analogue techniques are still being employed.

The Architectural Image

The modern world has played a key role in how we now define an architectural render, with modern technology changing our understanding of what an image can be and what it should convey. This is causing our definition of what an architectural drawing or render should look like to change from that of an architectural sketch to that of a rendered image. Architectural drawings have always been a powerful tool, the ability to manifest an image for something as grand as a skyscraper has helped to shape the way we build our communities today. Traditional architectural drawings were more akin to the work of a painter or artist, they were a subjective representation of a set of plans, and they were commonly recognized as just that, subjective. People understood that a hand drawn image of a planned building would be susceptible to an artist's interpretation and had the potential to change throughout its development. A completed project was not expected to look exactly like the initial drawings, because how was one person expected to be able to create all the intricacies and details of a complex structure by hand, on paper (**Figure 13**).

Figure 13: *A hand drawn perspective render from*
https://www.arch2o.com/from-stone-carvings-to-cad-how-architecture-drawing-has-evolved-over-the-years/

With the introduction of the photograph, the whole "landscape" started to change. Previously, a drawing was generally considered a subjective image, but with a photograph, we were now able to capture a frozen moment of reality (Llopis Verdú, 2018) and what was considered acceptable for an architectural image had just moved up a step on the ladder.

As photographs started to become more commonplace, what was expected from an image started to shift. No longer were people expecting an artist's rendition of a subject, but instead a snapshot of reality. At this time, both the hardware and software required for a computer to compete with a photograph was not even thought possible yet, so architectural images had to adapt in other ways. With the combination of photographs and technical drawings, architecture firms were able to piece together composites to represent proposed plans. While these could not directly compete with a photograph, they gave the firms an avenue of progression, to help further the architectural render.

The photograph and photograph-based montages and sketches would go on to dominate the architectural image space for a long time. Similar to architectural drawings, these images were not always representative of reality, often just showcasing the potential a building would be able to have under optimal conditions. "A painting or a drawing is always a codified artifice, a reflection on reality that is reinterpreted, recoded and transmitted, but it does not pretend to be reality" (Llopis Verdú, 2018) And just like previous images, this was mostly seen as acceptable, as creating an exact image of a building through sketching, and collaging existing photographs is a difficult task.

Figure 14: *An example of an architectural montage*
https://visualizingarchitecture.com/no-render-quick-collage/

As soon as technology was capable of rendering 3D graphics, development was started on creating software capable of rendering realistic scenes and images of not only buildings, but just about anything one can think of. These renders had a drastic impact on the way that perceive images and what photorealism meant. Before mainstream computer edited images, a photograph was representative. The legitimacy of an image did not need to be questioned. Now we have the ability to create photorealistic images that represent nothing but an idea, which has both positive and negative effects. Designers are able to create realistic renders that can accurately represent a planned project, however, due to the existing connotations that the general public has with photorealism, the photorealistic quality of these renders is often seen as fact. This is problematic when finished projects do not completely match the initial renders. Modern renders aspire to be the “definitive visual image” of a project rather than “provide conceptual information” about it (Llopis Verdú, 2018). With modern rendering software packages, we are now seeing hyper-realistic images used to represent planned builds, made to look a certain way that is unachievable in real life. This is making it harder for building companies to meet the standards clients and buyers have been led to expect.

Figure 15: *A modern hyperreal architectural render*
https://www.dezeen.com/2013/08/12/henry-goss-on-architectural-visualisations/

Legalities

Creating hyper-realistic renders is not inherently a problem, but when they are being used to deceive the general public, either intentionally or unintentionally, the work starts to delve into the realm of misinformation and false advertising. In New Zealand, the laws in place around this specific topic are vague and easy to circumnavigate. According to the Real Estate Authority (https://www.rea.govt.nz/, 2021), photos that are used to represent a property or building may be edited and altered, as long as the finished product is not depicting a scenario that would be infeasible to achieve with future renovations. Therefore, photos showcasing a residential household may be edited to show a wall removed, or different colour paint throughout the interior and exterior of the house, as these are seen as feasible changes that could occur in the future. However, removing or altering parts of the image that are out of the control of the owners is against New Zealand's real estate laws. Common examples of these are removing powerlines from an image or altering the background of a house to trick the viewers into thinking the property has more space or is not as close to other buildings compared to reality.

Where many images start to breach these rules is when lighting and foliage are adjusted to make the pictures more desirable. Similar to how fully computer-generated renders often run into issues with unintentional temporal conflation, doctored real estate photos often end up branching into location conflation. Real estate agencies that are editing photos by either replacing the foliage in exterior pictures or adding foliage where there was none before, are often using foliage from plants which could not be grown in that zone. For example, photos have been found with trees that are non-existent on the property and would take upwards of 40-50 years to grow to that size. Many photos also run into issues with unrealistic lighting, by enhancing the existing lights or creating artificial lighting scenarios, often resulting in rooms appearing to be naturally lit even though the sun would not naturally follow the shown path. While these more minor changes are not enough to warrant any form of legal action, they are still creating hyper-realistic images that are being used to deceive the audience and make the properties more desirable than what could be done without digital assistance. One of the more common photo manipulation tactics that does often encourage the authorities in New Zealand to step in and take action, is the editing or removal of any municipal structures that would normally been seen in the photography, such as powerlines, pylons and transformers. With the property owner generally having no impact on the placement and alterations of these, editing them out of photos or hiding them with digitally placed assets is a breach of New Zealand's real estate advertising laws and can see heavy cash fines imposed and the potential loss of a real estate license if caught.

False Advertising

While more traditional forms of hyperrealism often refer to paintings and other common artistic mediums, the vast majority of hyperrealism that the general public is exposed to is through the creative space that is advertising. Advertising dates back thousands of years, with examples of wall paintings and adverts on papyrus being found in ancient Greece and ancient Rome, around 5000 years ago (Lynn, 1958). Fast forward to the 19th century and the invention of the modern newspaper, companies now have a method to mass produce content and distribute it to the general public with relative ease and advertising really starts to gain traction as its own industry. Modern publications often refer to this specific time in advertising history as "the good old days, when a spade was a spade and people did not try to hide their weaknesses and frivolities, as they often do now" (Sampson, 2021). With the ability to reach the masses, advertising became much more rampant. As the space started to become more populated, companies started competing with each other, trying to make their ad stand out. Budgets for ads got bigger, as did the ads themselves. Which led us to scenes like the early 2000's New York Times Square, with larger-than-life billboards in every direction and companies trying to create larger and more eye-catching adverts than the last. Not long after that, mass adoption of social media and online advertising started to take over. With this space being even more congested than previous physical based advertising, companies had to really go above and beyond to try and set themselves above the others. As a direct result of this, more people opting for false or fraudulent methods to make themselves stand out from the rest of the crowd. As this process further developed, more cases of hyper-realistic images and false advertising have started to appear as companies needed to try and make their product really stand out.

Consumers are now in an age where media and advertising consumption is at an all-time high. In 2007, it was reported that the average consumer in the western world saw around 2,500 – 5,000 ads a day (Rhodes & Wilson, 2018). As of 2018, that number has increased to between 6,000 – 10,000 ads a day with the majority of those ads being seen online (Rhodes & Wilson, 2018).

High quality architectural renders are an excellent tool for smaller studios to have their work seen by their peers and the public and to make a name for themselves. However, not everyone is of the same opinion, with many considering the renders to be more of a farce that detracts from the architectural process. For example, Garcia Crespo (2011) said "According to many architects, their rendered images give a voice to constructive details, however, if those digital rendered images are only used at the end of the process, this is not architecture anymore, but advertising."

With the rise of social media and the ease at which a company or person can share information, posting renders and images online to spread awareness of a brand has never been easier. This ease at which information can be shared does come with disadvantages though. As a medium, the internet is a very unregulated space, with users preferring to self-regulate the content they are viewing and avoid government regulation getting involved (Ginosar, 2014). This often means that a company or person can advertise just about any content they want without their material having to be vetted. Because of this low level of moderation in the online advertising space, users are often subjected to large amounts of advertising imagery and often times not all of these images are legitimate, with only the user to determine whether or not the content they are viewing is trustworthy.

There is widespread evidence that false advertising is used by companies to overstate the value of their products (Rhodes & Wilson, 2018). However, the levels at which various vendors falsely represent their products online varies greatly, especially when talking in terms of hyperrealism. At one end of the spectrum, we see companies using images and videos that are drastically enhancing their content or are not related to the content they are selling. This is often seen in mobile advertising, especially games, as there are even fewer rules and regulations in that space. Using video content to lure in audiences with false claims and then keeping them there with predatory monetisation strategies, which currently falls under a marketing grey area but is proving to be an effective method with the younger generations (Graham, Young, & Marjan, 2021).

On the other end of the spectrum, there are companies using images and videos that do not accurately portray the product that they are selling. Many companies are using accurate source material but make subtle changes, resulting in hyper-realistic content that is often more appealing to the customer but not representative of the actual product. One of the more common examples of this is editing the colour pallet to the point that the colours would not be naturally occurring in the physical world (Struth, 2014). The vividness of an image and contrasting colours are important for how the audience perceives the product or an idea. (Struth, 2014). Struth (2014) talks about how the American coal industry was able to portray the idea of “Clean Coal” using hyper-realistic images of enhanced vibrant blue skies and emerald, green forests. These edited images were portrayed in parallel to the drab images of coal and coal related machinery, so that the American government was able to create positive and greenwashed connotations for many when it came to coal. These hyper-realistic false advertising campaigns made using relatively simple editing techniques, were able to convince millions of people that coal was a clean energy source.

Hyperrealism in other industries

As rendering technology has advanced and users become more skilful, the number of renders and images that are produced which are heavily edited and often hyper-realistic has dramatically increased, both at the amateur designer and the high-end developer level. These changes are often, especially to the trained eye, rather obvious. However, the architectural rendering industry is relatively new and has a far smaller consumer base when compared to other tech industries that are currently employing hyper-realistic techniques to help enhance their products.

One such industry is the cell phone camera industry. This industry has had decades of research and billions of dollars of investment and development put into creating the perfect end product image. There are just under 15 billion cell phones currently operating worldwide, and that number is expected to grow to nearly 20 billion by 2026 (Statista, 2022). Therefore, the demand to make an outstanding cell phone camera product is becoming increasingly important. One of the major ways, among others, that the various cell phone companies are doing this is with their camera hardware.

For the first decade, since the introduction of the camera in a cell phone, taking the same photo with two different brands of cell phones with similar specifications in their cameras would generally result in almost identical photos. However, as the competing companies have tried to differentiate themselves from their competition, cameras with similar specifications on different branded phones in 2021 often produce wildly different end results (Brownlee, 2022). Much of this is to do with the automatic processing that is happening behind the scenes with the camera software.

One of the first publicly documented cases of this happening was with the Nokia Lumia phones released by Microsoft in 2012 (Microsoft, 2022). When enabling the front facing camera and taking a photo, the phone would take five different photos in sequence and use a facial detection algorithm which looked at whether the subject's eyes were open and where they were looking, it would pick what it determined to be the best photo. As a result, the end photograph would not necessarily be a true representation of the photo that was initially taken. This, however, was not necessarily hyperrealism, as it was still choosing an image that the camera naturally took. Since the create of the Nokia Lumia, the use of software determined photo modifications have increased a great deal.

In 2019, Huawei released their new flagship phone, the P30 pro (HUAWEI, 2022). The P30 was considered by many consumers and reviewers to have the most advanced camera in a cell phone to date with a 5x optical zoom and up to 50x digital zoom and a 'moon mode' which automatically applied what the phone considered to be the best camera settings for capturing photos of the moon. Naturally, this led a lot of users posting close up photos of the moon (**Figure 16**). Immediately, people were suspicious and accusing Huawei of changing the photos during the post processing pass and replacing the moon with prestored photos (Jhaveri, 2019). Huawei released an official statement, claiming that this is not true, and the photos have been taken naturally with the cell phones. Further studies done by Zhang (2022) delved deeper into this topic and determined that Huawei's software was using artificial intelligence to enhance the photos of the moon. The software artificially brightens and darkens certain parts of the image and even goes as far as adding details that were not captured by the user. So, while the statement from Huawei was not wrong in saying their software was not superimposing existing photos into its users' photos, they were digitally altered so they looked completely different to the original photo and thus, created a hyper-realistic image.

Figure 16: *High fidelity photo taken on the moon with the P30 Pro (https://www.huaweicentral.com/week-one-whats-your-best-moon-mode-shot/)*

As the competition between cell phone companies rapidly increased over the years, the capability of the hardware within the phones developed equally as fast. The Pixel 6, the latest smartphone offered by Google at the time of writing, boasts both, a regular rear shooting camera as well an ultrawide camera. When taking a photo with the rear facing camera, the Pixel snaps a photo with both the standard and ultrawide lens, with the standard camera operating with a normal shutter speed and the ultrawide operating with a much quicker shutter. The camera software then detects whether or not the subject taken by the regular camera is blurry, if so, it then looks at the photo taken by the ultrawide camera with a quicker shutter speed. If it deems the face from this lens to be of better quality, it will superimpose the less blurry face onto the original blurry photo, resulting in higher quality, more recognisable image. Examples of this can be seen in the video by Brownlee (2021) where before and after images are compared.

It could be argued that these previous examples of the moon and subject alterations are enhancing digital media in a positive way. Although they are creating hyper-realistic images that are not exactly representing reality and doing so without the consent of the end user, these images are generally harmless and are not being used to deceive an audience or to sell a product. However, not all examples of modern hyperrealism are producing content that is helping to enhance the final product.

With the sudden boom in virtual meetings due to working from home as a result of the COVID lockdowns in 2020, companies have been looking at ways to stand out in this space. Towards the end of 2019, Apple started to work in partnership with FaceTime to develop a system to help create the illusion of maintained eye contact while on virtual calls (M. Wilson, 2019). Due to the user constantly having to look at their screen, which for the majority of people is located below their webcam, this results in losing the sense of eye contact that would happen if two or more people were meeting in person. Apple MacBook developers have been looking at fixing this problem by calculating the position of the user's webcam and digitally editing the eyes to give the illusion that the user is looking directly into the webcam, artificially creating a sense of eye contact again (Apple, 2022). However, this effect is resulting in rather lifeless eyes that are just staring directly into the webcam for the majority of the time. It takes rather exaggerated head and eye movement for the calculations to determine that your eyes are looking somewhere else and then disable the effect. This hyper-realistic attempt at creating a sense of eye contact is strongly bordering on being on the verge of the Uncanny Valley and lead to most users trying to figure out how to turn it off. A Google search of 'Facetime eye tracking' results in hundreds of web pages giving the users detailed instructions on how to disable the automatically enabled effect.

Greenwashing

Another example of hyperrealism and misleading advertising that is being seen seeing in today's world is, Greenwashing, a problem that is affecting many industries, not just those that are architecture related. Greenwashing is defined as "the act of misleading consumers regarding the environmental practices of a company or the environmental benefits of a product or service" (Delmas & Burbano, 2011). Often times, these misrepresentations are done using digitally edited photos, showing a product or company to be more 'green'/environmentally friendly than it really is. Companies that are greenwashing, to appeal to a more environmentally conscious consumer base, are undermining the efforts of other companies actively lowering their carbon footprints. This creates a sense of distrust in the consumer base using 'green' products. While this is a problem facing many industries with today's environment conscious consumer base, architecture and particularly architectural rendering is being heavily influenced by greenwashed projects (Figure 17).

Figure 17: *An example of a Greenwashed Render (https://www.re-thinkingthefuture.com/rtf-fresh-perspectives/a640-a-round-up-of-criticisms-of-photorealistic-architectural-renderings/)*

The leading cases of greenwashing in architecture are generally companies lying about their manufacturing processes to give off more of a sense of sustainability. They are claiming to use recycled products or claiming to use ethically sustained products. These examples of green washing do not fall under the category of hyperrealism, they are more along the lines of blatant false advertising. Examples of greenwashing hyperrealism can be found in companies who are trying to claim awareness of environmental problems through visual means to create the façade of being environmentally sustainable, a notable accolade that many new construction projects are trying to achieve.

Sustainability is a key consideration in modern architecture and can often be misinterpreted as greenwashing. Sustainability is often achieved in one of two ways in architecture; creating a structure using environmentally conscious and sustainable resources or creating a building that is considered carbon neutral. The former is often a much more expensive process than practices that do not take carbon emissions into consideration and provides extra challenges regarding structural integrity. The latter is often achieved through building practices that incorporate more natural lighting and ventilation, energy conscious designs and the inclusion of plants and foliage built into the design. For example, in Copenhagen, Denmark, city planners are trying to achieve a carbon neutral city by 2025 and one of the main driving factors that is helping to achieve this goal was the implementation of the ‘Green roof’ policy. In 2010 the city initiated a mandate for all new buildings being constructed with a roof slope of less than 30 degrees and all new roof renovations of existing flattop roofs to become green roofs, with the inclusion of plants and grass being planted on the roof (Donovan, 2015) (Figure 18). With the combined potential of over 200,000 square meters of available green roofing to be completed, the effect would result in quantitative progress to achieving carbon neutrality using a carbon credit system.

Figure 18: *Examples of the green roofing policy in Copenhagen (https://blueberriesconsulting.com/en/en-copenhague-ya-son-obligatorias-las-azoteas-verdes/)*

Projects like the Copenhagen green roof project require years of careful planning and have substantial positive outcomes for the environment. But for every legitimate example of an environmentally conscious architecture project, there is the potential for many more to be Greenwashing their plans. To properly design an energy efficient building requires an architect to have an in-depth knowledge across a range of different disciplines. Unfortunately, architects and architectural firms often lack this knowledge and fail to incorporate true sustainable architecture into their designs (Heine, 2013). Instead, these firms create hyper-realistic greenwashed renders (**Figure 17**). For example, they will showcase buildings of various sizes covered in trees and other foliage to convey the idea of sustainability in their projects. These designs will often create a structure and then add ‘green’ elements to it near the end of their pipeline as an afterthought, without consideration to the sheer amount of weight, space and logistics needed for a building to support that amount of greenery. As a result, the public is bombarded with these renders of impossible 'green' buildings, detracting and delegitimising from the work of companies who are legitimately incorporating sustainable design in a physically achievable manner.

Distribution Channels

The advancements in technology have been one of the major drivers in helping to progress the architectural rendering space. As the hardware and software has developed, the quality at which graphical engines are able to render images has greatly increased. Allowing digital artists to reach more viewers through multiple different mediums, such as social media platforms and utilising newer hardware like augmented and virtual reality. This allows for more exposure to architectural renders than previous more traditional distribution methods, such as the various paper-based outlets like magazines and newspapers.

Magazines

Magazines were once a powerful tool for distributing architectural images, for both newcomers and trendsetters alike. They were considered excellent material for new trends within a profession and were often an indicator for what could be the new style in in the years to come (Brown, 1968). But as times changed, and digital mediums started to be introduced, the way of the paper-based publication has started to lose traction, although paper-based media does still garner attention with the older population. However, with websites likes Instagram, a photo and video sharing social networking service, reportedly having over 2 billion current users (BusinessOfApps, 2022) the amount of people that an image or advertisement can reach is magnitudes higher than that of any paper-based publication that has ever been created. As previously mentioned, photography and architecture go hand in hand. So, when a social platform like Instagram comes along, which is purpose-built for displaying photography, it quickly becomes a much more attractive platform for architects to showcase their renders and ideas (Render, 2020).

Social Media

While Instagram is an excellent platform for architects to establish an online presence and promote themselves, it is also a platform that is heavily tailored to creating appealing images. Instagram not only encourages its users to edit their images beyond photoreal, but it also provides the tools for the users to do it with ease. Users can add filters to images, adjust exposure, lighting and edit other general camera postproduction settings. These tools help to facilitate the ongoing problems of unrealistic content creation on Instagram. The large number of users can be a negative for digital designers trying to show off their renders and work. When there are so many others posting pictures of their architectural renders, users start to run into the problems of a saturated market. Furthermore, if designers are not using the available filters and editing techniques available directly on Instagram, designers will often struggle to gain the same user engagement than those who do. This is resulting in a plethora of hyper-realistic architectural renders that are trying to be as eye catching as possible, just to garner more attention.

It is through the massive adoption of social media platforms like Instagram that more people are being exposed to architectural renders, both hyper-realistic and regular images. Previously, if someone was wanting to view architectural content, they would have to specifically seek it out. They would have to visit specific forums or websites or seek out a physical medium like books and magazines. People using these forums are generally more familiar with the subject material and will be able to determine when an image they are looking at has been edited or contains elements that would not be physically possible. These experienced specialists are able to filter information and not be influenced or tricked by exaggerated pictures. However, in large apps like Instagram, any user has the potential to be served an image with just about any subject material. This results in an audience who is unfamiliar with architectural renders consuming digital content that is often times heavily edited and unrealistic. It is scenarios like these where hyper-realistic content is having a negative impact on the community. With no previous exposure to architectural images, new viewers develop a warped perception of realistic architectural designs, resulting in legitimate work being dismissed or looked upon negatively due to it not being as eye catching as the hyper-realistic images.

Virtual Reality

Virtual Reality (VR) is a technology that allows a user to experience a computer-generated simulation of a three-dimensional image or environment, generally by wearing a head mounted display. In the early 1990’s, computers started possessing the imaging capabilities to move away from cartoon like images and towards more photo-realistic, convincing and accurate representations of physically based environments (Pittman, 1992). While at the time, these may have been considered high fidelity graphics, it is no longer the case as technology improved even further over the years (Hilfert & König, 2016). VR technology has advanced to the point where refresh rates are now high enough to instantly react to head movements without having to worry about the user experiencing sickness. While the hardware has developed, the software has also done so at a similar rate, allowing for true photo realistic visuals within the virtual environment, of significantly higher quality than previously possible.

While many people are starting to notice hyperreal aspects in standard computer edited photography (due to the increased awareness around the subject), the relative age and lack of public exposure to purpose-built VR environments can often make it harder to be aware of hyperrealism when viewing VR content. Due to the incredibly immersive nature of VR, a user can often have difficulty discerning irregularities when viewing a scene in VR, meaning that hyper-realistic elements can often slip by unnoticed.

VR is a powerful tool when it comes to displaying photorealistic architectural content. It offers many interdisciplinary opportunities for developers to design and create buildings in an interactive environment (Portman, Natapov, & Fisher-Gewirtzman, 2015). However, as the technology becomes more accessible and more architectural content is made with VR, it will have to be monitored to see the effects that hyper-realistic scenes have on the display and advertising of architecture.

While VR can be a very effective tool when it comes to having users experience 3D environments, it is not without its shortcomings. Due to a user being required to wear an expensive headset, it is not feasible to market the experience to the average person. Often the only way to demonstrate a scene in VR is to have the user come to you. This can make VR as a medium rather impractical and severely limits the audience. However, as cell phones are becoming more powerful, it possible to place them specially designed enclosure that can be worn on the head and create a VR experience (Samsung, 2022).

Augmented Reality

Augmented Reality (AR) is a technology that superimposes computer-generated imagery onto a user's view of the real world, creating a composite of the two. As AR applications are becoming more popular and viable, more industries are starting to incorporate them. There is an upwards trend of AR utilization in practical fields such as education, design, manufacturing, architecture and more (Chi, Kang, & Wang, 2013). The nature of AR with the overlapping of digital assets augmented over real-life footage, allowing the user to easily alter and add to reality, this lends itself perfectly into hyperrealism. With one of the only constraints being that cell phone hardware is not as powerful as PC-based rendering systems, users are able to create superimposed scenes with few constraints and minimal thought being put into realism. Because AR only requires users to have a cell phone or similar hardware, augmented environments have a much greater audience reach than VR.

With the overlaying of digital objects over real world footage, it is much easier to add to an augmented reality environment than it is to subtract from it. The ability to continuously add to an environment results in scenes with many hyperreal elements, however, unlike VR and other entirely digital mediums, the mix of digital models and real world models can hamper the perceived realism of scene because of the noticeable difference between the digital and the real (Kim, Kim, Borrmann, & Kang, 2018). Which in turn can detract from any hyper-realistic notions that were created in the AR environment.

How 3D rendering is influencing architectural rendering.

The recent progress made to 3D rendering software packages is one of the more impactful developments to happen to the architecture industry and architectural rendering practices since their inception. However, the ease and access to create these hyper and photorealistic renders is not seen as a good thing by everybody. Depending on the profession, there are both positive and negative comments about the impact that 3D rendering is having.

Architectural renders are known to be changing people's perception on the architectural industry, and for the purist, that can be considered a negative thing. Some are viewing architecture as more of a fashion statement and a commodity and the rendered image has been lambasted for perpetuating the notion of style over substance and the image over the experience (Ross Bryant, 2013). Some of the larger, more prestigious architectural firms almost look down on computer renders, holding a preference for the more 'traditional and artistic' way of creating a render. Some argue "Resorting to this type of realistic image is to a certain extent debatable when used in contexts where more understandable and operative systems of representation exist, namely architecture colleges or even design competitions." (García Crespo, 2021). However, many smaller studios are unable to just ignore creating hyper-realistic and attractive 3D renders as they do not have the established connections to operate, and these renders are a good way of showing the public what your firm can create. Browsing through architecture related message boards and forums, it is easy to find a plethora of articles, blogs, videos and forum posts making valid points for both sides of the argument. For example, some detailed, well laid out blog posts go into elaborate detail into how computer visualisations are impacting Jim Keen's architectural practice and those around him (Keen, 2022). Keen argues that as computer renderings have become more commonplace, "they have started to become a liability more than a tool" when it comes to expressing an idea. The main purpose of architectural renders are to communicate an idea, Keen argues three main points in which he thinks the popularity of hyper-realistic renders is negating that.

Firstly, Keen (2022) believes hyper-realistic computer visualisations are overused. As previously discussed, with the current accessibility of the various modelling and rendering software that exists today, it has never been easier for one to learn how to create a visually stunning render. This, in Keen's opinion is leading to an oversaturation of hyper-realistic renders, meaning that they have lost their impact factor.

Secondly, community users are obsessed with creating the perfect render, so much so that the amount of work that goes into creating a perfect render (for a competitive market) is starting to detract from the bigger picture, which is creating renders that display the functionality and purpose of a design. Focusing on such small details, such as making sure that shadows are calculated correctly, or that the decorations on the coffee table do not look too fabricated, results in both the designer and the audience often losing sight of the bigger picture. These small objects are derailing the narrative and causing consumers to forget about the overall design and impact that the render is supposed to be imposing. This focus on perfection also gets in the way of connecting with less experienced audiences. Often times, these renders are created at the concept level. Regardless of how much detail and work has gone into them, they are generally just an idea or an interpretation of proposed plans. Yet architects often get held to and are represented by these renders. A displeased audience will dismiss the render because there are elements they do not like, without realising that these are subject to change. While on the other end, people will complain that certain elements of the render did not end up making it into the completed building, also not realising that the render was not an exact representation.

The last point that Keen (2022) argues is that of monotony. As the various modelling programs advanced over the years, most of them are now capable of producing high quality renders which are quite similar to each other. This means that presentations are starting to become quite homogenised, often containing a mix of illustrator diagrams and high-quality renders. Keen (2022) argues that for designers to stand out from the rest of the crowd, they should steer away from photorealistic rendering and try to develop a unique art style that places emphasis on the message you are trying to portray with your content (**Figure 19**).

Figure 19: *An example of stylized rendering (https://architizer.com/blog/practice/tools/computer-renderings-old-news/)*

While Keen (2022) does raise some good points around the topic of homogenisation and overuse, it sounds like he is tackling the issues from a rather elitist point of view that might not apply to a large amount of the architectural visualisation community, due to the size and number of renders that smaller firms put out. Others, including Geilman (2018) and Easyrender (2020), have opposing views to Keen (2022), advocating for the continual growth and mass adoption of rendering packages and advancements to the field.

Geilman (2018) and EasyRender (2020) go into detail about various new technologies and how they are helping to advance the field of architectural visualisation. Geilman (2018) talks about how photo realistic renderings allow for designers and architects to create life-like imitations and more accurate representations of their final product. These advancements make it possible to capture an idea more precisely and perfect a design before it is constructed.

EasyRender (2020) share a similar view but take it a step further, emphasizing how 3D rendering software such as Rhino, Sketchup and Maya are revolutionizing how users work in the digital world by manufacturing opportunities to create and develop like never before. With the inclusion of publicly available software such as Google Earth, users are now able to access high detailed terrain and satellite imagery, further aiding in how users can plan and develop building propositions.

As Geilman (2018) points out, architecture itself is benefitting from these technological advancements. It is allowing users to view buildings in a way that was not previously possible, giving architects more insight into how their creations will function in the real world. The quick turnaround on realistic renders gives them the ability to more easily determine how an idea functions with realistic lighting, objects and people in the space.

Hyperrealism is becoming more prevalent in the architecture rendering field. The steady rise in scholarly articles and uptick in hyperrealism related posts on architecture blogs and websites is confirmation of this. As the technology increases, the ease at which one can create hyperreal imagery has increased by magnitudes in the last 10 years. This is allowing enthusiasts to produce work that was previously only created by high-end architecture firms and those with an extensive, in-depth knowledge on how the software works and the mathematical principles behind it. This increase of accessibility and the ability to easily share work through social media and other online platforms means that not only is more hyper-realistic content being made, but it is also much easier to view said content. Based on current trends and the data available, it is safe to assume that as the technology continues to improve, there will continue to be a rise in hyperreal architectural content being produced.

Chapter 4. How to make a hyperreal image

In this section, I am going to break down some of the main technical aspects that go into creating a realistic 3D render. This information will be useful for people looking to get into both architectural and general 3D rendering, as well as for people who are already familiar with rendering, but wanting to make their renders more realistic.

I will break this down into four key sections (as previously discussed in Chapter 2, The technology behind the render) that have a significant impact on the fidelity of the final output including lighting, texturing, modelling and camera work. While there are many aspects and nuances to these four topics, I will cover the ones that I find to be the most impactful from my perspective.

I will produce various renders for each of these four key elements. The software packages used to create the various 3D elements are Maya for 3D modelling (Autodesk, 2022), Photoshop and Substance Designer for texture creation and texture editing (Adobe, 2022a) (Adobe, 2022b) and Redshift for lighting and rendering (Maxon, 2022).

Lighting

Firstly, I will take a look at lighting and how it can impact the look and feel of a render. With all of these renders, I will try to keep the scene as simple as possible so that other aspects do not influence the overall aesthetic too much.

Lighting a render is very contextual and will vary greatly depending on different variables. For example, if the scene is indoors or outdoors, or if are one is trying to replicate studio lighting or natural lighting. Lighting has two main roles, both have a great impact on the final output of a render. The main purpose of lighting a rendered image is to make the objects visible in a scene. While that sounds relatively simple at first, there is more to it than just placing a light in the middle of a scene. Lighting can be used to direct the viewers' attention, to hide and reveal information or to help set a certain mood or tone. The second element of lighting is the creation of shadows. This is often overlooked, but accurate shadows go a long way in making a scene look photoreal. Shadows are used to help perceive depth in an image and to give objects a sense of being grounded. Poorly created shadows are a quick way to break immersion and are often the deciding factor in when determining whether an image is rendered or not.

Figure 20 depicts a simple untextured chair with a single direction light.

Figure 20: *A single directional light is a quick method of lighting but produces harsh shadows.*

Many artists will create outdoor scenes with a single directional light to replicate the sun. This results in dark shadows and very hard shadow lines. Lighting a scene like this creates a very large amount of shadow-based contrast and causes a lot of information to be lost in shadow. If someone was to photograph a house or object with deep shadows like this, it would generally be considered a poor photograph, the same will generally be said about a render with the same characteristics.

Figure 21 uses the same directional light as **Figure 20** but also contains a fill light.

***Figure 21**: A single directional light with fill lighting results in much softer shadows and a better lit chair.*

The fill light creates a softer shadow and eliminates some of the hard, self-shadowing created by the geometry of the chair. The softer shadows are more akin to a real-life scene where light is bouncing off various surfaces, helping to illuminate the shadowed areas. These fill lights are generally required in both indoor and outdoor scenes to generate a more realistic light. There is no set rule as to how many fill lights one needs and where to place them. This will vary scene to scene and depend on the content depicted.

When lighting scenes with multiple light sources (**Figure 22**), generally found with interior scenes, there should be visible contact shadows with objects touching, or low to the ground.

Figure 22: *The chair lit with multiple light sources, most commonly found in interior scenes.*

These will become less obvious as one moves away from the light sources and should match the strength of each light source in the scene. One will generally want to have a main shadow that is being cast from either the sunlight or a hero light, the strongest light in the scene. The other light sources should have an impact on existing shadows in the scene, making them much duller and giving them a softer appearance. Having multiple light sources in a scene is commonly used to fill in dark areas left by stronger directional lights. This results in a well-lit scene with shadows that are not overly dark and obscuring objects.

When it comes to effectively lighting a scene, what happens off-screen can have just as much of an impact as what is in the frame. There are two main factors that are heavily impacted by off-screen objects and materials, and this is highlighted in both **Figure 23** and

Figure 24. **Figure 23** shows how having objects out of frame can impact the global illumination calculations that are going on and the impact that can have on a scene.

Figure 23: *Off-screen global illumination is being used to cast colour over the scene.*

Using the same colour and intensity lighting as **Figure 21**, but with the addition of blue and red objects placed at differing distances off-screen, one can see a tangible effect on the outcome of the lighting. When the red object is closer to the right of the frame (as in **Figure 23**), there is a very apparent red light bouncing back into the scene. This is dulling the appearance of the shadows while also casting a red hue about the majority of the scene. With a smaller blue object out of frame to the left, there is a more subtle blue glow being cast. The effect these out-of-frame objects can have on a scene are often more impactful with indoor scenes as there are more objects and walls for light to bounce off compared to an open outdoor space.

Having various coloured bounce lighting in a scene is key when it comes to creating realistic renders. Real world lighting scenarios, unless in a controlled space, will have many different objects reflecting different coloured lights around a space. Without this, scenes can start to look washed out as the scene will be flooded with a single colour of light.

Another factor that is often overlooked with populating the space outside of the frame is the fidelity of the reflective surfaces that are visible on screen. While it is possible to generate high quality reflections using a High Dynamic Range Image (HDRI) map, to create accurate and ‘true’ reflections, users need to have a fully populated space for those reflective surfaces to bounce light off.

Objects out of frame are often forgotten about when considering how the shadows in a scene should look. This applies to both indoor and outdoor scenes. Off-screen shadows help to set the tone of a scene and should be more apparent when the sun is at a lower angle (**Figure 24**).

Figure 24: *Off-screen shadows have the ability to add character to a scene.*

They can be used to help install a sense of scale and create boundaries without having them visually defined in frame. Some of the most common examples of off-screen shadowing generally stems from foliage, especially in outdoor scenes, and various house framing when looking at indoors scenes.

Texturing

Like lighting, the texturing of various objects (and scenes overall) can be quite a subjective practice. There are times where the designer will want to create a one-to-one 3D representation of a real object, and other times they will want to add their own take to environment, either to elicit a certain feeling or to portray a desired aesthetic. Through texturing, artists are able to do this.

As discussed in the section 'The Technology behind the render' on page 25, texturing with a Physically Based Rendering (PBR) workflow involves four main texture maps: Albedo map, Roughness map, Metallic map, Normal map. There are other maps that can be used in the texturing process which can control various elements such an ambient occlusion and emissive textures, however, they are often considered as extras and not necessary to create a PBR material. Each map is used to control and convey a certain aspect of information to the viewer. Understanding how these different aspect work individually and in conjunction with each other is an integral part of the 3D rendering pipeline.

Albedo map

An Albedo map, often also referred to as a Colour map is used to control the colour of an object. This map will be the basis of the entire material, so working with a good Albedo map is essential when creating realistic materials (Denham, 2022). An Albedo map will usually be made up with either a flat colour or a texture file depicting a more complicated pattern of colours, such as bricks or wood. When working with a detailed Albedo map, such as a brick or wood pattern, it is best practice to try and work with an image that has flat lighting. Using an image that has shadows already baked into it will either limit the lighting to coming from a certain direction as to match the baked in shadows or will cause problems when trying to adjust the lighting as there will be shadows baked into the texture as well as the shadows generated by the lights. **Figure 25** is an example of a wood Albedo map with and without both shadows and lighting baked in.

***Figure 25:** An example of an albedo map with shadows and lighting baked in.*
(https://www.artstation.com/blogs/luismesquita/jGXd/everything-about-pbr-textures-and-a-little-more-part-2)

While Figure 25A looks more realistic than Figure 25B, it is very limiting in terms of creative control. With the shadows and lighting highlights already baked in, the lighting will need to be set up exactly to match the existing shadow data. Any other lighting configuration will result in “double shadows” and incorrect colour information.

Normal map

A Normal map is used to generate the illusion of depth on an otherwise flat surface. However, these are usually limited to smaller scale bumps and indents. To properly generate bumps and indents on a larger scale, it is best to either model them in initially or to use a bump map, which I will touch on later.

Every polygon inside of any of the engines used for rendering 3D images has what is called a surface normal. These surface normals are used to control how those surfaces react when light hits them. In a rather simple sense, the angle between a polygon's surface normal and the direction of in the incoming light source can be used to calculate how lit or shadowed the polygon's surface should appear. The colour information in the Normal map is used to modify the direction of the surface normal for every pixel and therefore can achieve an illusion of fine detail and depth. Due to the way the information is encoded in the RGB colours, a normal map will generally have a pale blue appearance. A more detailed breakdown explaining the mathematics behind surface normals and normal maps can be found in LearnOpenGL (2022).

Figure 26 is a comparison of the same surface without and with a Normal map, respectively. The Normal map adds a perceived depth and a large amount of detail to the surface when compared to the surface containing no Normal map. More importantly, this detail will change significantly when the direction of the light is changing.

Figure 26: *A surface comparison without and with a Normal map.*

These flat surfaces are a common problem that can result in a viewer being able to differentiate between a rendered image and a photograph. Real life surfaces often contain many signs of wear and teach such as scratches and dents. The inclusion of a Normal map containing this information, such as shown in **Figure 27**, helps to eliminate the clean, mathematically produced look that 3D engines create by default.

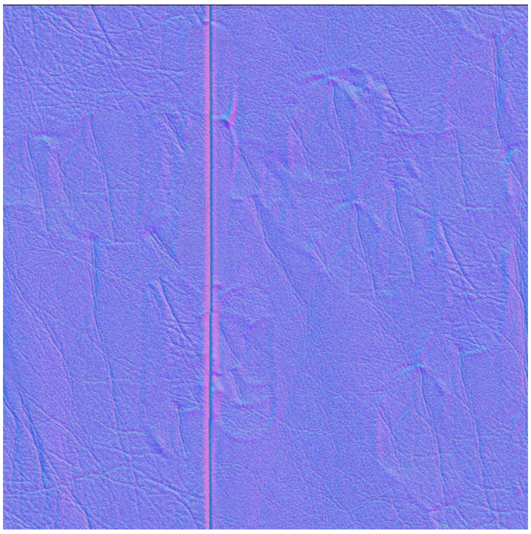

Figure 27: *The Normal map used to add extra detail in* ***Figure 26***

Roughness map

Roughness maps, also sometimes referred to as Glossiness maps, are used to control the physical roughness of a surface. They do this by defining how the light scatters when it hits the surface of an object. Roughness maps are a greyscale map that operate with values between 0 and 1. A value of 0, which corresponds to black on the Roughness map will cause that area of the surface to not scatter any light. As a result, the surface will reflect all the light hitting it back outwards. Visually, this would be akin to a mirror, or a very shiny surface as shown in **Error! Reference source not found.**. At the other end, a Roughness map will have a value of 1 and will be white. This means that light hitting the surface will be completely scattered over the surface and will have a fully diffused look.

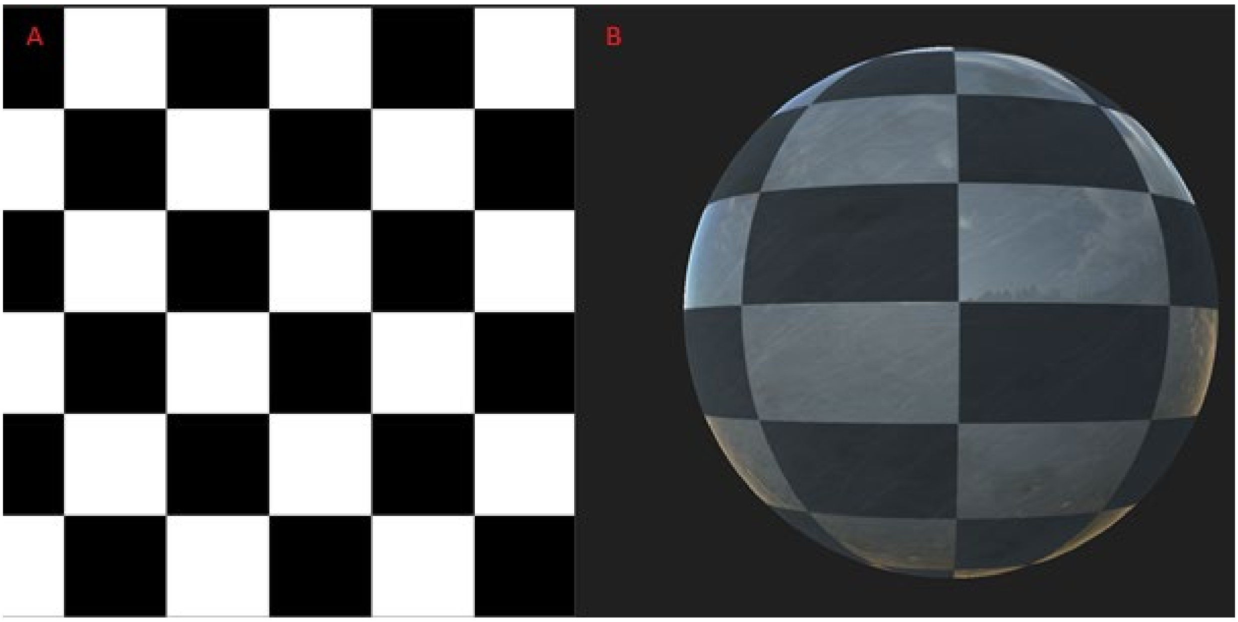

Figure 28: *An example of a Roughness map and the effect it has on a surface.*

Roughness maps differ from Metallic maps in the sense that they still inherit information from the Albedo map. One can see in **Error! Reference source not found.**B that the reflective patches have a black base, which is based on the Albedo map. If the albedo was instead red, the reflective surface would have a red hue to it.

The Roughness map can be considered one of the more creative aspects of the PBR maps as it gives the artist the opportunity to visually define the character of a surface. The roughness values can be used to tell a story about a surface and to better define what sort of environment it has been in. In conjunction with the Normal map, the artist is able to define whether a surface has been looked after or treated roughly, or whether or not a surface has been exposed to the elements or if it has been sheltered. Roughness values can be quite subjective and getting creative with them can add significant detail to a render.

A common mistake made, particularly in architectural rendering is both undervaluing and overvaluing roughness values. Most real-world surfaces will have some form of roughness value. There are very few surfaces that would need a roughness value of 1. Scenes lacking an adequate amount of roughness will look very flat and two-dimensional as lighting will not be bouncing as much and there will be a severe lack of reflections. Another common mistake is artists placing too much emphasis on the Roughness maps in their scene. While some may choose to do this as an artistic choice, it makes the surfaces in a scene look wet as all the objects are shiny and reflecting too much light. A good balance in roughness values helps to create the contrast between the surfaces and generally follows best practice guidelines.

Metallic map

Of the four main base maps, the metallic map sees the least use due to the way it functions. Operating with the same greyscale maps as the Roughness map, Metallic maps are used to control what parts of an object are metal and which are not. A value of 1, or a white image, represents a fully metallic surface and a value of 0, or a black image, represents a surface with no metallic properties. The metallic map works in conjunction with the albedo. A value of 0 will mean that the Metallic map has no effect on that surface and the colour will be represented by its respective map. In contrast, a value of 1 means that the Metallic map will completely override the Albedo map and replace it with a metallic look. Any value in between will give the user a mixture of the Albedo map and some level of a metal surface, depending on the values as seen in **Figure 29**.

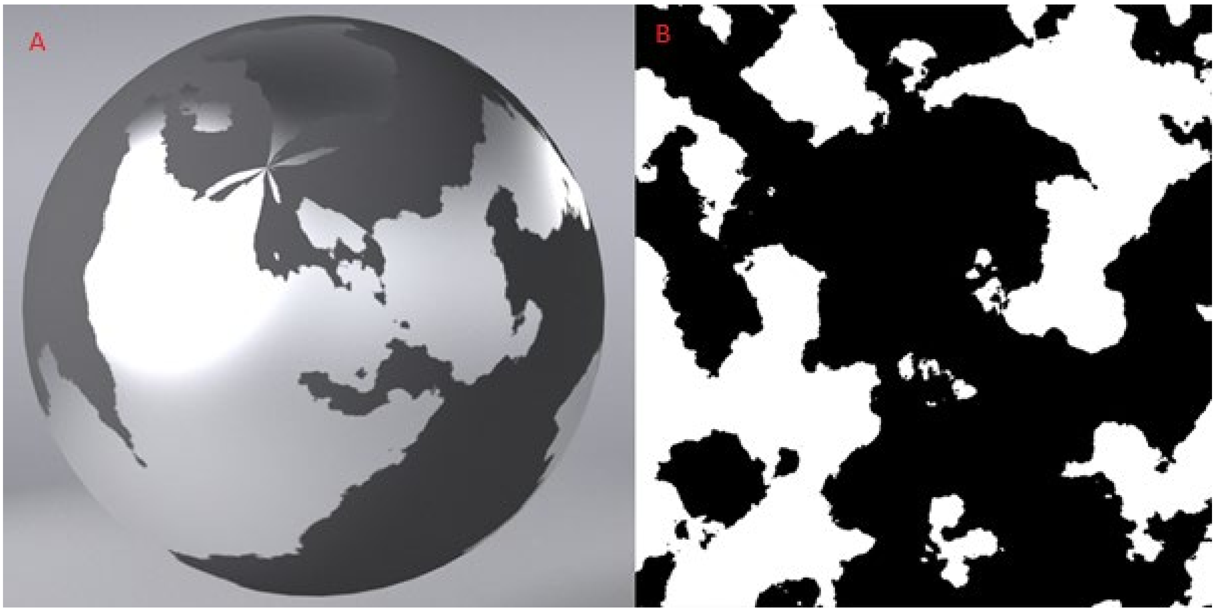

Figure 29: An example of a metallic map and the effects it has on a surface.

As previously mentioned with the Roughness map, Metallic maps do not inherit any of the colour information from the Albedo map. As seen in **Figure 29**A, the metallic patches of the sphere have not inherited the darker colour of the Albedo map, instead resulting in a silverish metallic surface. If desired, a user can control the appearance of a metallic surface with the use of a Specular map. This map can help give the appearance of coloured metals or be used to better reflect back coloured lights.

With the proper use of the Albedo, Roughness, Normal and Metallic map, an artist is able to create a material that will look physically accurate regardless of the scene that it is placed in. However, there are various other maps that are not required for a PBR material to function but enhance or add to the workflow, including the Ambient Occlusion map, Height map and Emissive map.

Ambient Occlusion map

The Ambient Occlusion (AO) map is used to define how much ambient bounce lighting is accessible to designated area of a surface. A proper AO map should result in more shadowed areas on various surfaces where light has more trouble reaching. This occurs where two surfaces meet and create a slightly more shadowed area. Similar to baking shadows into the Albedo map, baking the AO map into the Albedo will result in inaccurate shadows in certain lighting conditions.

Height map

Height maps are often confused with the Normal map. While they perform quite similar functions, they do so in very different ways. Where the Normal map creates the illusion of depth and shadowing via the surface normals, the Height map actually adds geometry to the mesh through a principle called tessellation. This can be used to enhance the effect of a Normal map or to generate completely new geometry. The Height map also works off a greyscale image. Black represents the baseline level of the mesh and white will create the highest peaks. Shades of grey will represent values in between.

However, tessellation is a very computational heavy practice exercise and in the majority of cases, modelling the extra geometry during the 3D modelling stage will produce a better result. One advantage is that the use of a height map does allow the user to adjust the height of the geometry on the fly through scaling the values of the height map. With the current state of the hardware available to users, tessellation is not recommended when trying to develop in VR.

Emissive map

The Emissive map is a colour map that is used to indicate surface areas that should emit instead of reflecting light. The colour of the light is based on the respective colour of the map. These materials are particularly handy when trying to replicate LED displays and strip lighting. While entire scenes can be lit using emissive materials, it requires far more work than traditional lighting techniques. Overuse of emissive lights can easily result in an environment looking washed out and lacking detail as well as being more computationally difficult to render.

Modelling

Of the three main components to making a realistic representation of an object, 3D modelling is the least subjective of the three. The most realistic 3D model is generally the one that resembles the real-life object the most (Ilias, 2017). This leaves very little room for creative freedom when it comes to architectural modelling. If the designer is trying to achieve a stylized look, then there are far more creative choices that can be made, allowing for a more subjective modelling experience.

When trying to create a realistic 3D model, the designer has the potential to both undershoot and overshoot in terms of realism. With models that take on a more organic shape, such as objects that are not man-made or have surfaces with more shape and variety to them, there is more of a chance of these being under-realistic. The more detailed surfaces require a much more knowledgeable and talented designer to be able to recreate them accurately. However, this is more of a problem in other fields than it is in architectural rendering, due to the less organic shape of most architecture. Models that are more geometric by design quite often fall into the category of being over-realistic, with overly straight lines and perfect angles. The models start to take on the appearance of an object that is computer calculated and made, which is much more of a problem in architectural design.

When modelling for architectural scenes, a conscious effort needs to be made to avoid an overuse of exact 90-degree angles. By default, all 3D modelling applications will create a box with perfectly sharp 90-degree corners. This results in what would be technically impossible edges in a real-world scenario. Visually, this causes two major problems, an excessive amount of symmetry and unrealistic, hard edges. These edges do not pose too much of a problem when working with the exterior of architectural projects as the camera usually is not placed close enough to a surface to notice the hard edges. However, when visualising interior environments and placing the camera closer to objects, these infinitely hard edges, as seen in **Error! Reference source not found.** start to give objects a very boxy look and become an easy way to differentiate between a render and real photograph.

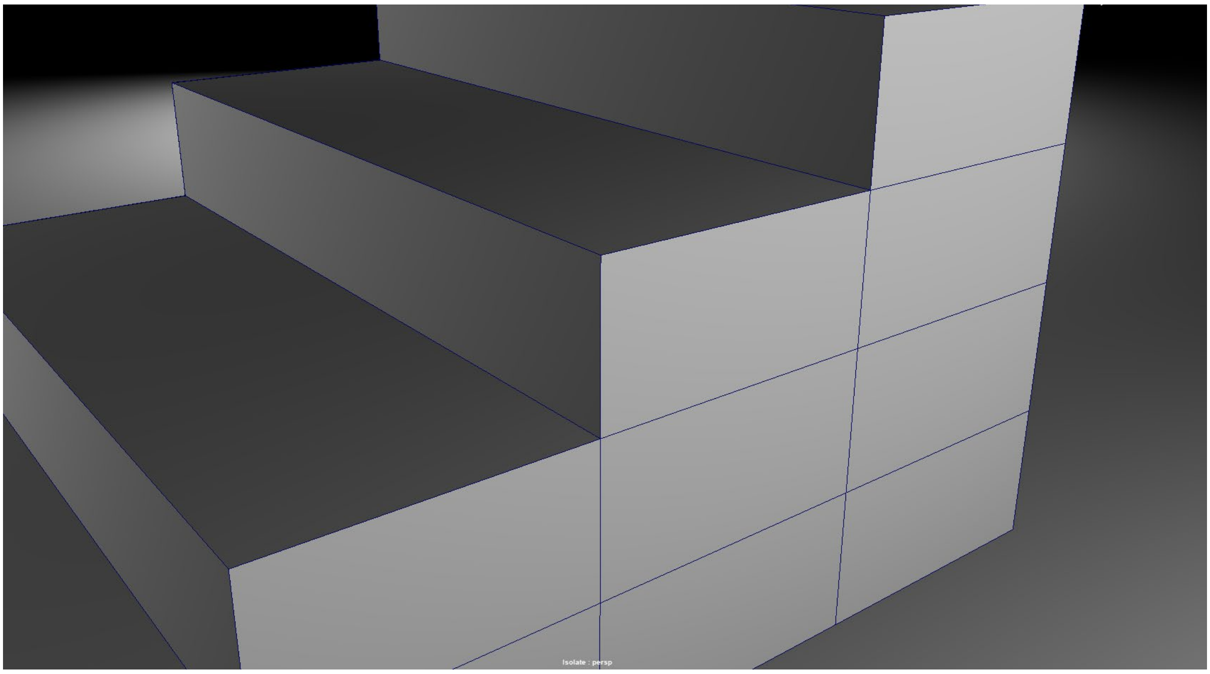

Figure 30: *A model of stairs created with perfectly square right angle geometry.*

There are two common methods to fix hard edges in most modelling software packages. The simpler of the two is to double edge right angles. As seen in **Error! Reference source not found.** adding a second edge is a simple way to artificially create rounded geometry on corner edges. While effective, adding a second edge to every hard surface can be a labour-intensive method that does not give the designer much control. Real world carpentry and architecture often used what is called a chamfer, the equivalent of adding a second edge to geometry to create a 45-degree face as opposed to a hard 90-degree corner.

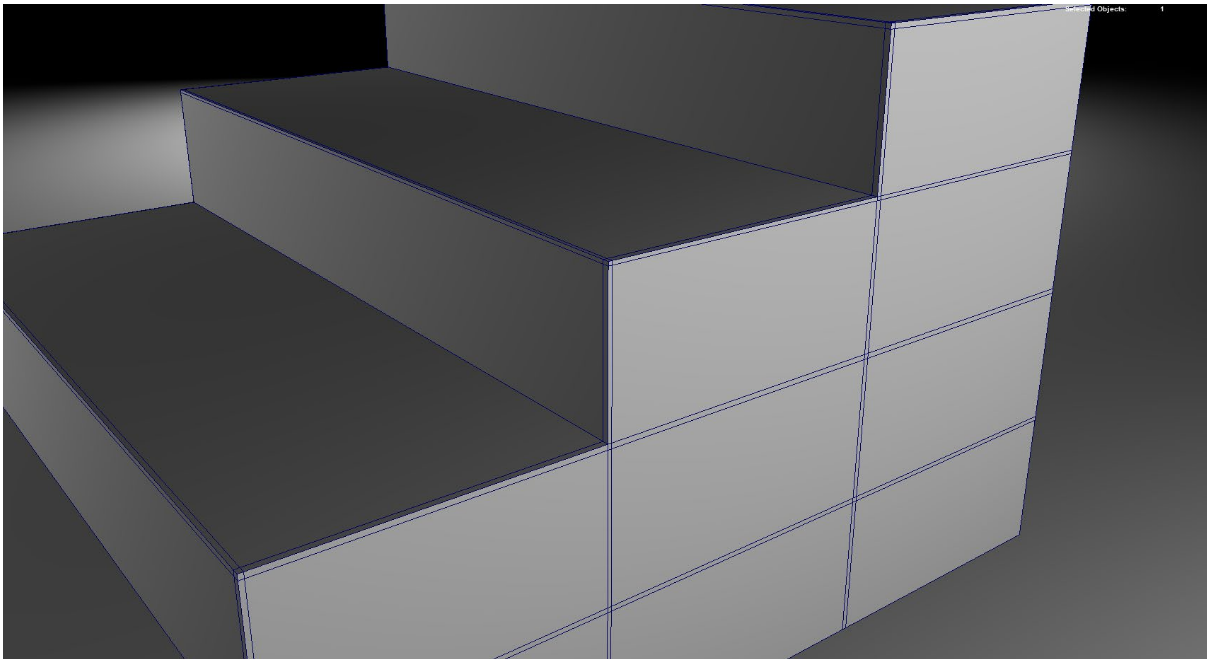

Figure 31: *An example of a model using double edged geometry to help give the appearance of smoother edges.*

The second method is the use of the 'Bevel' tool. More commonly used than just double edging models, the Bevel tool gives a user far more control over softening the edges. By adding a bevel, the user can control how many edges they want to add and the distance between them, giving them the option to create a more subtle bevel all the way to a completely rounded edge as seen in **Error! Reference source not found.**. The bevel tool often gives the best results, but it comes with the downside adding the most geometry to the scene of the three different options.

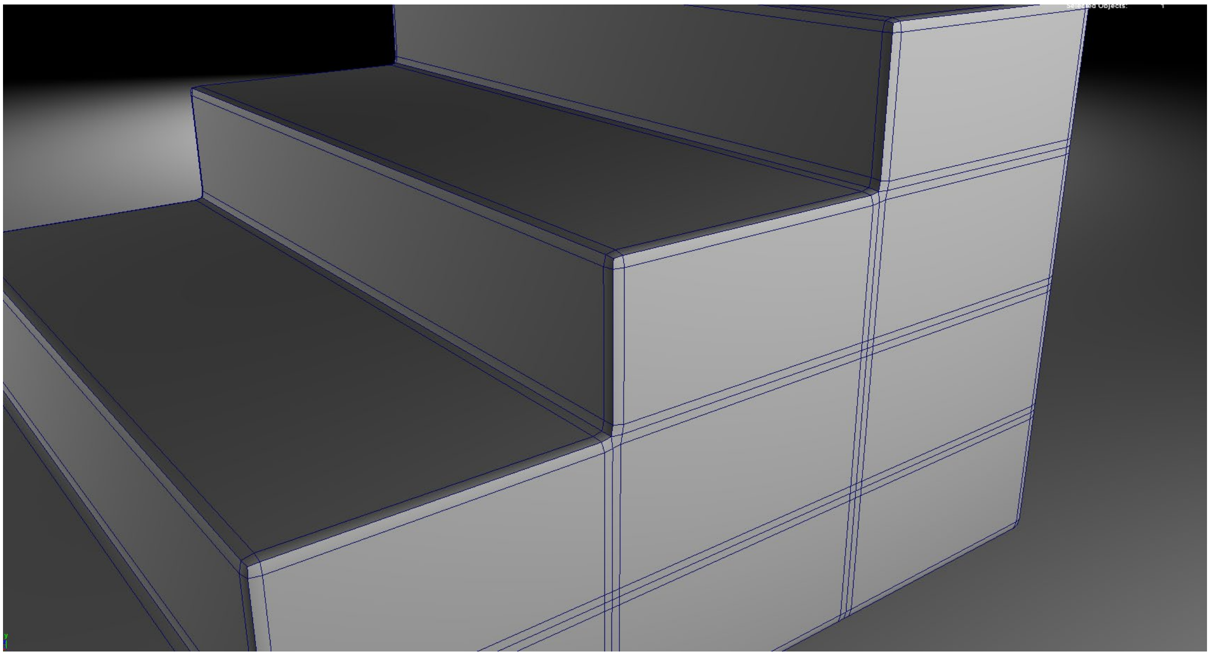

Figure 32: An example of a model created using bevelled edges to smooth the hard edges of the geometry.

The overwhelming number of symmetrical objects is something that effects both exterior and interior renders. While for some, this symmetric look may be a stylized choice that the designer is trying to achieve, for others, it gives off the aesthetic of a computer simulated scene. Breaking up overly symmetric shapes and patterns by modelling in small defects or moving objects around to break the symmetry can help to make a scene feel more human.

Camera

Depth of Field

There are many aspects of a digital camera that can be altered which have a great impact of the visual fidelity of a render. Often, these settings are passed over and the realism of a scene suffers for it. As discussed in the section 'New Vs Old' on page 6, the closer a render looks to a photograph, the easier it is to convince the average person that it is real. Adding a sense of depth to a scene is a relatively easy way to achieve this, as shown in Figures 33 and 34. **Figure 34** is a render with depth of field enabled, creating an obvious foreground and background. In contrast, **Figure 33** is the same image without any depth of field, leaving the entire scene in focus. Adding depth of field helps to direct the eye and gives the image much less of a computer-generated feel. In this case, the added depth of field is used to help hide minor details that are incorrect with the models, when an image is fully in focus it is much easier to spot these blemishes. Depth of field will vary depending on the scene, especially when comparing interior and exterior designs. Interior shots will often use a narrower depth of field with a more aggressive focal length which results in a more aggressive blur (**Figure 34**).

Figure 33: *An example of an interior render without depth of field.*

Figure 34: *An example of an interior render with depth of field enabled.*

For exterior shots, a wider depth of field is often wanted as this is used to place more of an emphasis on both the foreground and midground. This allows the designer to keep an entire building in focus while framing it with foreground objects that do not draw the viewers eye away from the subject matter (**Figure 35**).

Figure 35: *A render with a blurred foreground (sourced from https://www.artstation.com/artwork/6azl5W)*

Perspective

Perspective is one of the most important camera functions when it comes to replicating photorealistic imagery. With the ability to place a digital camera anywhere in a three-dimensional space, it is quite easy to use the camera in a way which would be physically impossible in the real world. Until recently, with the introduction and rise in drone photography, most photography that people were exposed to was taken from the ground level.

Drastically changing perspectives does a couple of things. First, it changes the eye line of a photo which changes the vanishing point (the point where hard lines meet in the distance), which are prominent in objects with hard, symmetric lines, such as buildings. A warped perspective on the vanishing point makes it harder for humans to perceive things such as depth and scale. Secondly, it exposes viewers to an angle that they are not accustomed to.

Placing a camera in a virtual scene that is looking down on a house creates a perspective that is unfamiliar due to a lack of photography from that angle. For many, this can be enough to break the illusion of a hyperreal image and to determine it is instead a render.

Resolution

By default, rendering engines set the resolution to a 16:9 aspect ratio. This is the current standard at which computer monitors are produced at (Steam, 2022), therefore it is the default resolution that renders are often created at. Changing the resolution to something other than 16:9 can give a render a more unique feeling and help to differentiate it from the standard 16:9 digital content such as **Figure 35**. Similar to how people have acclimated to the photograph and what that can represent, people are starting to become accustomed to the 16:9 digital image. Moving away from that can help a render stand out when all the other content being produced is the same resolution.

While there are more elements that go into creating realistic 3D renders than those that have been listed, these cover the fundamentals required to produce a virtual environment. When taking all these different elements into consideration, it is possible to create images with both photorealistic and hyper-realistic elements such as seen in **Figure 36** which uses lights that could not physically exist to enhance the lighting and shadows.

Figure 36 : *An example of a scene put together with photorealistic and hyper-realistic lighting elements to achieve nicer shadows.*

Chapter 5. Analysis and Discussion

Hyperrealism is a relatively new concept in the architectural rendering community, slowly gaining traction over the last 10 years. The advancements to rendering packages are allowing architecture studios, both large and small to create renders to a much higher fidelity than what was previously possible. With an awareness of hyperreal images and the effect that they can have on an audience, it is possible to visit existing architecture firm's websites to see if they are using hyperreal images to advertise any of their previous or existing projects. Based on whether or not a company is advertising with hyperreal images and the types of projects they are being used for, we can gain insights into when and how firms may use hyperreal images that might inadvertently misrepresent or overemphasize their projects. In order to explore this concept, I will analyse architecture firms from Auckland, New Zealand by going through their online portfolios and examining how they advertise the various projects they have worked on.

Firms were excluded from this analysis if they had less than five employees, and if they did not have at least one registered architect. Without a registered architect, they would otherwise be ineligible to make a bid on government grant projects which is where a large amount of architectural rendering appears to be done. While this does still leave many firms, a focus will be placed on those that rank higher in google searches and have larger portfolios to work with.

Architecture Firms

After applying the search limitations and adding an arbitrary limit of 20, the resulting architectural firms that will be examined are listed as follows, in alphabetical order.

Archimedia Group Ltd
Architectus New Zealand
Athfield Architects
Cheshire Architects
Chow:Hill Architects
Context Architects
Ignite Architects
Isthmus
Jasmax
Leuschke Group Architects
Monk MacKenzie
Oxygen Architecture

Patterson Associates Limited

Peddlethorp

RTA Studio

Stevens Lawson Architects Ltd

TOA Architects

Warren and Mahoney Architects

Wingate Architects Limited

WSP Auckland

These listed firms all have at least five employees with multiple registered architects and advertise an extensive range of projects, both completed and currently ongoing (as of November 2022), through their website portfolios. Within this group, five firms have been selected that offer a wide range of exhibited images which will be presented and compared.

Patterson Associates

At first glance, there is a good mix of both rendered images and photography that this company uses to promote their practices. Patterson Associates (Pattersons, 2022) is a moderately sized company, with a team of architects and technicians that fluctuates around 30 - 40 people. They appear to be making the conscious decision to promote their work using photography from previous projects on the front page of their website (**Figure 37**).

Figure 37: *Michael Hill Club House, from the front page of Patterson Associates (https://pattersons.com/commercial/michael-hill-golf-clubhouse/)*

However, this does not mean that they are not using hyperreal renders to represent their work elsewhere in their portfolio. When looking further into their work, you can see them presenting projects in both the ‘Civic’ and ‘Residential’ space using hyperreal renders such as **Figure 38** and **Figure 39**.

Figure 38: *A hyperreal residential render created by Patterson Associates (https://pattersons.com/residential/singhania-house/)*

With the lighting in **Figure 38** being inconsistent with the sun placement in the sky to ensure that each part of the scenery is lit to provide maximum exposure, the team at Patterson’s are using images with unrealistic lighting simulations to promote their architectural concepts. **Figure 38** is also an example of a firm 'Greenwashing' an image, with a heavy focus on unrealistic foliage to make a building look 'cleaner' and more desirable. **Figure 39** is another example of a hyperreal render that can be found after looking through the projects on their website. This render was used to enter the hotel tender competition in 2013 and was the winning design selected for the creation of the Park Hyatt Hotel.

Figure 39: *A hyperreal 'Civic' render created by Patterson Associates (https://pattersons.com/civic/park-hyatt-hotel/)*

When compared to the now built Park Hyatt Hotel seen in **Figure 40**, it provides a clear example of how a hyperreal render will not necessarily be representative of the finished product. In **Figure 39**, the team at Patterson's have chosen to omit the surrounding docks and other buildings as well as any other structures that are present in the surrounding skyline. The heavily edited sky only serves to add to the unrealistic setting of the building as well as the reflections of the non-existent, distant skyscrapers seen in the reflections of the exterior walls. Patterson's, like many other firms are making the conscious decision to remove the immediate surroundings out of their renders to make their design more attractive.

Figure 40: *The finished building for the Park Hyatt Hotel (https://auckland-hotels.co.nz/wp-content/uploads/2020/12/Park-Hyatt-Auckland-5-Star-Waterfront-Luxury-Hotel.jpg)*

Patterson's also have examples of hand drawn proposals, montage and computer-generated renders such as **Figure 41** that do not fall into hyperreal territory. However, most of these images were created at a much earlier date, with some even going back to the late 1990's. This could be an indicator that the advancements in the technology are what have allowed this practice to start creating and distributing hyperreal content. Whether it was intentional or not, it is plain to see that over time Patterson's have started to use more images with hyperreal elements.

Figure 41: Watercolour image created for a waterfront village proposal (https://pattersons.com/civic/orakei-waterfront-village-2010/)

Cheshire Architects and RTA Studio

Many of the architecture firms listed seem to promote their work in a similar way to Patterson's, with a good mix of photography and renders, both hyperreal/photorealistic and stylised. However, there are a couple of firms on that list who prefer to promote their practices almost solely with photography, such as Cheshire Architects (Cheshire, 2022) and RTA Studio (RTA, 2022). These two firms appear to differ in the clientele that they are trying to attract to their respective businesses, although they are both opting to use similar methods to do so.

Cheshire Architects are a medium size studio who employ around 25 people (as of November 2022) and were originally founded in the late 1980's. They are a bespoke architectural firm who advertise their work almost exclusively through photography such as **Figure 42**.

Figure 42: *A photo promoting 'Waiheke House' made by Cheshire (https://www.cheshirearchitects.com/residential/cowes-bay-cheshire)*

Cheshire promote themselves with an even spread between residential, commercial and cultural projects. However, unlike other firms, it does not appear that they make bids for public tender projects, and looking at renders such as **Figure 39**, this is the space that a large amount of the hyperreal renders are being created in. With the goal of making your firms designs stand out above the rest and look visually stunning, incorporating hyperreal elements becomes much more inviting.

Based on their website, the way they promote themselves and the photography they choose to display (**Figure 43**), Cheshire Architects could be considered more of an artisanal, personalised architecture firm that has more of an interest in residential architecture and much less of a focus on commercial and public projects. It could be that Cheshire consider architectural renders to be less artisanal or less traditional than high-end photography and do not want to diminish their brand. Or it could be their focus on residential architecture. But either way, it seems that they opt not to advertise their brand with either hyperreal renders, or architectural rendering in general.

***Figure 43:** High-end photography of 'The Folly' project by Cheshire Architects (https://www.cheshirearchitects.com/residential/the-folly-cheshire)*

It is worth noting that due to the large amount of post processing and the way that Cheshire Architects choose to frame their photography, quite a few of their photos start to take on some of the traits that would often be seen in a render, as seen in **Figure 44**. With the way this image is framed, the slightly more exposed interior and such a wide depth of field, it would not be surprising for many to think this was computer generated without a more in-depth inspection.

Figure 44: An example of a photo looking like a render. (https://www.cheshirearchitects.com/residential/eyrie)

RTA Studio, which was founded in 1999 are a slightly smaller studio, employing anywhere between 10-15 people. Overall, they have received more than 100 international and New Zealand architectural awards. Including, New Zealand Home of the Year, three World Architecture festival category wins and the New Zealand Architecture Medal (RTA, 2022). They, like Cheshire Architects, place an emphasis on photography to advertise their work as seen in **Figure 45**.

Figure 45: *Pollen Street Office, designed by RTA Studio (https://rtastudio.co.nz/portfolio/pollen-street-office)*

However, unlike Cheshire, RTA Studio does a considerable amount of work in the public space, including commercial spaces and public building such as schools and libraries. When looking through RTA's online portfolio, all images that are representing both the residential and public projects are photographs of existing works done by the studio. It is only when looking at their commercial projects that there are renders being used (**Figure 46** and **Figure 47**).

Figure 46: *Ice Hotel proposed plans in Queenstown, NZ. (https://rtastudio.co.nz/portfolio/ice-hotel)*

Figure 47: *Te Wero bridge proposal by RTA Studio (https://rtastudio.co.nz/portfolio/te-wero-bridge)*

RTA Studio's online portfolio consists of over 60 projects. Of those projects, only two contain renders; the Ice Hotel plans in Queenstown and the Te Wero bridge proposal in the Auckland Viaduct. Both of these projects are proposals that were not accepted but made it into the final round for contention. So, unlike other practices who often showcase renders on their portfolios to represent buildings that have been completed, RTA Studio is only doing so to showcase their unfinished proposals that they feel are good enough to advertise their practice. Both contain hyperreal elements in an attempt to make their proposals more appealing. The Te Wero bridge proposal images make drastic changes to the colour and appearance of the water, while the Ice Hotel proposal makes heavy adjustments to the lighting and bloom effects.

RTA Studio and Cheshire Architects are both companies that place an emphasis on using photography to showcase their work. This could be because they feel that architectural renders detract from the authenticity of their work, and they wish to portray themselves as higher end firms. Or it could be because these are well established firms, and they have enough examples of previous completed projects that they can accurately portray their practices through photos of those existing projects. However, they are both well-established firms that have a certain level of prestige to them that are happy to showcase 3D renders.

Monk MacKenzie and Warren and Mahoney Architects

When looking through the list of architecture firms, two are clear leaders in the amount of renders they have available on their respective websites including Monk MacKenzie and Warren and Mahoney Architects. However, the two firms do differ considerably in some respects.

Monk MacKenzie ("Monk Mackenzie Architects," 2022) is an Auckland based firm that was founded in late 2013 and employ around 35 staff. They, unlike the previous mentioned firms, choose to promote most of their portfolio with computer-generated renders. When looking at the front page of their gallery, a user is presented with nine projects that their team have worked on. Out of the nine projects, six of these are rendered images, with varying levels of realism and heavy amounts of post-production work.

One of these projects is the proposed alpine site located on the Cardrona Valley floor shown in **Figure 48** and **Figure 49**. This series of photos appears to use a mixture of CG buildings and terrain/fauna superimposed over existing landscapes. This results in quite an uncanny image as the viewer is left trying to determine which elements are generated and which are real. Coupled with the questionable lighting and heavy film grain that is applied to not only this project, but many others, this render starts to border on hyperrealism.

Figure 48: *Proposed images of the alpine village in Cardrona (https://www.monkmackenzie.com/project/central-otago-master-plan)*

Figure 49: *The proposed alpine village located in Cardrona. (https://www.monkmackenzie.com/project/central-otago-master-plan)*

While Monk MacKenzie choose to represent themselves mostly through rendered images, many of these renders are of a photorealistic quality and would be quite hard for the average person to differentiate between photo and render, such as **Figure 50** and **Figure 51**. Considering that Monk MacKenzie is a relatively new company of nine years (at the time of writing), it could be they do not have enough completed physical projects to promote their portfolio. It could also be, that as a firm, Monk MacKenzie has the skills and technical knowledge to create computer-generated imagery that is indiscernible from photography, for the average consumer, and are happy to use those images to showcase their work.

Figure 50: *Photorealistic render created by Monk and MacKenzie (https://www.monkmackenzie.com/project/arahia)*

Figure 51: *Photorealistic render created by Monk and MacKenzie (https://www.monkmackenzie.com/project/edition)*

Warren and Mahoney (Mahoney, 2022) are one of the oldest, still operating architecture firms in New Zealand. They were founded in 1955 and currently employ over 150 staff. They have a broad selection of previous projects to put on display to represent their company as they have been in existence for nearly 70 years, yet they still opt to display a large amount of rendered imagery, with a significantly varying level of realism between images, such as **Figure 52**.

Figure 52: *An example of rendered imagery from Warren and Mahoney.*
(https://warrenandmahoney.com/portfolio/gardeners-road)

Figure 52 is an example of a render that is not exactly hyperreal, but just heavily edited to the point where it is obvious that it is a render, causing it to lose its photorealistic attributes. Out of the listed architecture firms, Warren and Mahoney appear to have the most variability when it comes to the quality of renders that they display on their website, not only from project to project, which can be somewhat expected from such a large firm, but even from image to image when promoting the same project, such as in **Figure 53** and **Figure 54**. Both images are listed on their website for the reconstruction of the Christchurch cathedral. However, they are of vastly different quality with **Figure 54** trying to make up for this fact with hyper realistic lighting and large amount of bloom.

Figure 53: *A high quality render of the Christchurch cathedral by Warren and Mahoney.*
(https://warrenandmahoney.com/portfolio/christ-church-cathedral)

Figure 54: *A lower quality render of the Christchurch cathedral by Warren and Mahoney.*
(https://warrenandmahoney.com/portfolio/christ-church-cathedral)

Considering the number of architects that Warren and Mahoney employ, it is easy to see why the majority of their work is done on larger scale projects in the civic and public sector as there is considerably more money to be made than in private housing. Working in the public sector generally tends to lead towards more digitally created content, as architecture firms often need to compete with others to be accepted to work on the project. This means that renders are used as hypothetical plans to try and secure a project. As we have seen with other architecture firms, most are making use of the newer technologies afforded to them to create renders are close to photorealistic as possible, with some incorporating hyper-realistic elements, either on accident or by design. While there are some examples of Warren and Mahoney trying to achieve that same standard (**Figure 54**), many of their projects are still being promoted using less than photoreal renders, with many of these projects having been from 2020 and onwards, such as **Figure 55**.

Figure 55: *An example of the lower quality renders Warren and Mahoney still produces.*
(https://warrenandmahoney.com/portfolio/northcote-aquatic-recreation-centre)

Not only are their renders of a quality that does not meet the standards of today's computer-generated imagery, especially for a firm as large as Warren and Mahoney's, but there are also renders of questionable quality that are including hyperreal elements. Projects such as the Pescara Bridge design (by Warren and Mahoneys) use images that appear to contain a mix of 3D rendered imagery and photo montage, as well as the inclusion of hyper-realistic lighting and fog. Resulting in an image that is quite uncanny and does not do a good job of representing the project (**Figure 56**).

Figure 56: *One of the images used for the Pescara Bridge design by Warren and Mahoney.*
(https://warrenandmahoney.sgp1.digitaloceanspaces.com/images/7352_Pescara-Bridge_2.jpg)

As previously mentioned, Warren and Mahoney have a deep history of projects that can be used to represent their portfolio, and besides some of the questionable renders they use, they do also have high quality photographs (**Figure 57**) and photoreal renders on display (**Figure 58**).

Figure 57: *An example of photography used to represent Warren and Mahoney's portfolio.*
(https://warrenandmahoney.com/portfolio/waihinga-martinborough)

Figure 58: *An example of photorealistic renders that Warren and Mahoney use.*
(https://warrenandmahoney.com/portfolio/sera-brighton)

When looking through the online portfolios of the listed architecture firms, there is a clear discrepancy in the way the varying practices advertise their works and the type of aesthetic they are trying to promote. There are firms, like Cheshire Architects, who are aiming to advertise towards more of a niche market, offering a bespoke service and doing so almost exclusively through photographic material. This decision by Cheshire could imply that they find renders to be less artisanal or that they think renders could cheapen their brand. It is worth to note that Cheshire Architects mainly work in higher end residential architecture. Other firms are more than happy to represent their work with an assorted range of renders. However, they are mostly doing so when showcasing buildings and designs in the public sector as these are often used in an attempt to secure the design of the building. Firms of all sizes are often using hyperreal renders in their portfolios, whether they do so on purpose or not is unknown. To the layman, these renders may appear to be good representations of work, but to the more experienced eye, the hyperreal elements often detract from the original design and are a net negative on the overall image.

Future Work

Future study in this space could involve looking at more fields than just architecture. Any field that uses computer-generated imagery to promote or to sell their services or products could potentially stand to gain from using hyperreal elements in their work. A more in-depth conversation with existing firms about when they are choosing to create CGI to promote their work and how they are choosing to style it would be beneficial in future studies. For example, to understand whether the architecture firms are aware of hyperrealism when creating work with hyper-realistic elements in it, or if hyper-realistic products are being produced largely due to technological advances that have occurred in recent years.

Furthermore, it would be interesting to conduct a poll or survey on the public using various imagery, with a mixture of photographs and rendered images, as well as images both with and without hyperreal elements to see if there is any bias towards of the listed options.

Chapter 6. Conclusion

Hyperrealism is a relatively new movement that it is becoming more popular as people are able to produce it more easily as new technologies develop and become available. Despite its sudden rise in usage, it is fair to say that most consumers are unaware of hyperrealism. For those that are aware of its existence, they may question the efficacy of its usage in various visual fields, such as architectural visualisations.

Hyperrealism is an effective tool that has the ability to make a designer's work stand out from others in a space like online advertising, which is quickly becoming an oversaturated market. However, like Jim Keen (2022) said, when everyone starts to make hyperreal renders, the effectiveness of them will start to lessen, and more unique styles of displaying work will start to become more prominent. This was explored in the portfolios of Auckland architectural firms in this thesis.

From looking at the listed architecture firms based in Auckland, New Zealand, it is clear to see that there is a disparity in the use of hyperrealism between both projects and firms. It is much more prevalent in the larger scale projects, such as public bids and townhouse developments, whereas smaller projects, such as residential buildings and smaller office spaces tend to favour more traditional photography over renders. This is most likely due to both monetary factors (with larger projects having larger funding pools) and the amount of people that will see a project. Often, the amount of time and skill required to generate hyper-realistic renders will not be financially viable when the audience is small and does not need to be convinced to invest in a development or project.

Hyperrealism is a unique art movement in the sense that it has the capacity to reach a much wider audience than that of other artistic movements due to its use in advertising and other public facing industries. Through mediums such as architectural visualisation, the layman is consuming large amounts hyperreal content, whether they are aware of it or not. Other artistic movements generally do not reach the public eye, for one to view them, they often need to be aware of the movement and voluntarily seek it out through spaces such as art exhibitions. This alone means that hyperrealism is able to deceive more people than other art movements, through either intentional or accidental deception.

As technology keeps progressing, hyperrealism will start to become more commonplace, and not just in architectural rendering. The ability to promote a product with images and videos that make a product look better than it would in real world scenarios is too valuable for marketing purposes and as it stands currently, too valuable in making your work stand out from the rest.

References

Adobe. (2022a). Adobe Photoshop. Retrieved from https://www.adobe.com/nz/products/photoshop.html

Adobe. (2022b). Substance Designer. Retrieved from https://www.adobe.com/nz/products/substance3d-designer.html

Amoakoh, S. (2020). What is Hyperrealism? Retrieved from https://useum.org/hyperrealism/what-is-hyperrealism

Apple. (2022). MacBook. Retrieved from https://www.apple.com/nz/mac/

Argan, G. C., & Robb, N. A. (1946). The architecture of brunelleschi and the origins of perspective theory in the fifteenth century. *Journal of the Warburg and Courtauld Institutes, 9*(1), 96-121.

Autodesk. (2022). Maya. Retrieved from https://www.autodesk.co.nz/products/maya/overview

Baianu, I. (2012). On Asymmetry in Biology and Nature. *Nature Precedings*, 1-1.

Bauwens, S., Bartholomeus, H., Calders, K., & Lejeune, P. (2016). Forest inventory with terrestrial LiDAR: A comparison of static and hand-held mobile laser scanning. *Forests, 7*(6), 127.

Blinn, J. F. (1977). *Models of light reflection for computer synthesized pictures.* Paper presented at the Proceedings of the 4th annual conference on Computer graphics and interactive techniques.

Borden, I. (2007). Imaging architecture: The uses of photography in the practice of architectural history. *The Journal of Architecture, 12*(1), 57-77.

Borden, I., & Rendell, J. (2000). *Intersections: architectural histories and critical theories*: Psychology Press.

Brown, D. S. (1968). Little Magazines in Architecture and Urbaimism. *Journal of the American Institute of Planners, 34*(4), 223-233.

Brownlee, M. (2022). The Blind Smartphone Camera Test 2021. Retrieved from https://www.youtube.com/watch?v=dp4nWm59esI

Bryant, R. (2013). "The addition of real-world imperfections is taking architectural visualisation to the next level". Retrieved from https://www.dezeen.com/2013/08/12/henry-goss-on-architectural-visualisations/

Bryant, R. (2013). Peter Guthrie Interview. Retrieved from https://www.dezeen.com/2013/10/20/peter-guthrie-on-hyper-realistic-visualisations/

Buchanan, K. (2011). The Biggest Problem With the Tintin Movie Might Be Tintin Himself. Retrieved from https://www.vulture.com/2011/07/the_biggest_problem_with_the_t.html

BusinessOfApps. (2022). Instagram Revenue and Usage Statistics (2022). Retrieved from https://www.businessofapps.com/data/instagram-statistics/

Chen, Y., Cui, Z., & Hao, L. (2019). Virtual reality in lighting research: Comparing physical and virtual lighting environments. *Lighting Research & Technology, 51*(6), 820-837.

Cheshire. (2022). Cheshire Architects. Retrieved from https://www.cheshirearchitects.com/

Chi, H.-L., Kang, S.-C., & Wang, X. (2013). Research trends and opportunities of augmented reality applications in architecture, engineering, and construction. *Automation in Construction, 33*, 116-122.

Clarke, G. (1997). *The photograph*: Oxford University Press, USA.

Conte, P. (2018). Hyper-Realism.

Debevec, P. E., Taylor, C. J., & Malik, J. (1996). *Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach.* Paper presented at the Proceedings of the 23rd annual conference on Computer graphics and interactive techniques.

Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California management review, 54*(1), 64-87.
Denham, T. (2022). Texture Maps: The Ultimate Guide For 3D Artists. Retrieved from https://conceptartempire.com/texture-maps/
Designblendz. (2019). The evolution of architectural visualization. Retrieved from https://www.designblendz.com/blog/the-evolution-of-architectural-visualization
Dixit, I., Kennedy, S., Piemontesi, J., Kennedy, B., & Krebs, C. (2019). Which tool is best: 3D scanning or photogrammetry–it depends on the task. In *Biomedical visualisation* (pp. 107-119): Springer.
Donovan, E. (2015). Architecture Washed in Green. *SUSTAINABILITY: An imperative for plurality and context*, 22-23.
Dostal, C., & Yamafune, K. (2018). Photogrammetric texture mapping: A method for increasing the Fidelity of 3D models of cultural heritage materials. *Journal of Archaeological Science: Reports, 18*, 430-436.
EasyRender. (2020). Influences of technology on architecture and design. Retrieved from https://easyrender.com/a/influences-of-technology-on-architecture-and-design
Eye, G. (2021). The Importance of Lighting in 3D Modelling. Retrieved from https://glancingeye.com/the-importance-of-3d-lighting/?lang=en
Fasse, I., & Paul, J. C. (1995). Realistic Rendering and Computer-Aided Lighting Design in Architecture.
Ferwerda, J. A. (2003). *Three varieties of realism in computer graphics.* Paper presented at the Human vision and electronic imaging viii.
García Crespo, J. M. (2021). High fidelity. The immature use of digital images for architecture. *EGA Expresión Gráfica Arquitectónica, 26*(42), 156-167.
Gavrila, D. M. (1999). The visual analysis of human movement: A survey. *Computer vision and image understanding, 73*(1), 82-98.
Geilman, G. (2018). How Technology is Revolutionizing the Architecture World. Retrieved from https://www.boxx.com/blog/emerging-technologies/how-technology-is-revolutionizing-the-architecture-world
Ginosar, A. (2014). Self-Regulation of Online Advertising: A Lesson From a Failure. *Policy & internet, 6*(3), 296-314.
Gosh, S. (1981). History of photogrammetry. *Laval University, Canada*.
Graham, C., Young, F., & Marjan, A. (2021). The generation Z audience for in-app advertising. *Journal of Indian Business Research*.
Green, J., Matthews, S., & Turanli, T. (2002). Underwater archaeological surveying using PhotoModeler, VirtualMapper: different applications for different problems. *International Journal of Nautical Archaeology, 31*(2), 283-292.
Greenberg, D. P., Torrance, K. E., Shirley, P., Arvo, J., Lafortune, E., Ferwerda, J. A., . . . Foo, S.-C. (1997). *A framework for realistic image synthesis.* Paper presented at the Proceedings of the 24th annual conference on Computer graphics and interactive techniques.
Hanson, D., Olney, A., Prilliman, S., Mathews, E., Zielke, M., Hammons, D., . . . Stephanou, H. (2005). *Upending the uncanny valley.* Paper presented at the AAAI.
Heine, U. (2013). Teaching sustainability in design without greenwashing. *David Publishing Company www. davidpublishing. com*, 395.
Hilfert, T., & König, M. (2016). Low-cost virtual reality environment for engineering and construction. *Visualization in Engineering, 4*(1), 1-18.
https://www.rea.govt.nz/. (2021). Real Estate Authority. Retrieved from https://www.rea.govt.nz/
HUAWEI. (2022). Huawei P30 Pro. Retrieved from https://consumer.huawei.com/nz/phones/p30-pro/
Ilias, A. K. (2017). What does the most realistic 3D model look like? Retrieved from https://www.quora.com/What-does-the-most-realistic-3D-model-look-like

Jhaveri, A. (2019). Are The Moon Shots From The Huawei P30 Pro Fake? Retrieved from https://in.mashable.com/tech/3114/are-the-moon-shots-from-the-huawei-p30-pro-fake
Keen, J. (2022). The Problem With Photo-Realistic Architectural Renderings. Retrieved from https://architizer.com/blog/practice/tools/computer-renderings-old-news/
Keskeys, P. (2022). 8 Ways to Fool Your Audience with Hyperrealistic Renderings. Retrieved from https://architizer.com/blog/practice/details/look-twice-hyperrealistic-architectural-renders/
Kim, H. S., Kim, S.-K., Borrmann, A., & Kang, L. S. (2018). Improvement of realism of 4D objects using augmented reality objects and actual images of a construction site. *KSCE Journal of Civil Engineering, 22*, 2735-2746.
Kullmann, K. (2014). Hyper-realism and loose-reality: The limitations of digital realism and alternative principles in landscape design visualization. *Journal of Landscape Architecture, 9*(3), 20-31.
Kunz, S. (2015). The problem of realism in animated characters: has the Uncanny Valley been crossed?
Lafortune, E. (1996). Mathematical models and Monte Carlo algorithms for physically based rendering. *Department of Computer Science, Faculty of Engineering, Katholieke Universiteit Leuven, 20*, 74-79.
Lange, E. (2001). The limits of realism: perceptions of virtual landscapes. *Landscape and urban planning, 54*(1-4), 163-182.
LearnOpenGL. (2022). Advanced Lighting - Normal Mapping. Retrieved from https://learnopengl.com/Advanced-Lighting/Normal-Mapping
Lin, T.-H., Lan, C.-C., Wang, C.-H., & Chen, C.-H. (2014). *Study on realistic texture mapping for 3D models.* Paper presented at the 2014 International Conference on Information Science, Electronics and Electrical Engineering.
Liu, G., Ceylan, D., Yumer, E., Yang, J., & Lien, J.-M. (2017). *Material editing using a physically based rendering network.* Paper presented at the Proceedings of the IEEE International Conference on Computer Vision.
Llopis Verdú, J. (2018). El paradigma fotográfico del dibujo arquitectónico digital. *Arte individuo y sociedad, 30*(3), 557-573.
Lvovsky, A. I. (2013). Fresnel equations. *Encyclopedia of Optical Engineering, 27*, 1-6.
Lynn, R. A. (1958). The Story of Advertising. By James Playsted Wood. New York, The Ronald Press, 1958. *Business History Review, 32*(4), 462-464.
Maganga, M. (2021). Architectural Rendering and the Slippery Slope of the Uncanny Valley. Retrieved from https://www.archdaily.com/959563/architectural-rendering-and-the-slippery-slope-of-the-uncanny-valley
Mahoney, W. a. (2022). Warren and Mahoney Architects. Retrieved from https://warrenandmahoney.com/
Manovich, L. (1995). The paradoxes of digital photography. *Photography after photography*, 57-65.
Markovic, K. (2018). *Authenticity in digital surrogates. Workflow development for generating an authentic digital surrogate for heritage conservation.* Auckland University of Technology,
Maxon. (2022). Redshift. Retrieved from https://www.maxon.net/en/redshift
Microsoft. (2022). Lumia Product Page. Retrieved from https://www.microsoft.com/en-us/mobile/phone/lumia640
Mitchell, W. J. (1994). *The reconfigured eye: Visual truth in the post-photographic era*.
Monk Mackenzie Architects. (2022). Retrieved from https://www.monkmackenzie.com/
Moreira, S. (2021). Who Are We Making Hyper-Realistic Renders For? Retrieved from https://www.archdaily.com/960849/who-are-we-making-hyper-realistic-renders-for
Mori, M. (1970). The uncanny valley: the original essay by Masahiro Mori. *IEEE Spectrum*.
Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine, 19*(2), 98-100.

Morrison, A., & Skjulstad, S. (2010). Unreal estate: Digital design and mediation in marketing urban residency. In *Exploring Digital Design* (pp. 189-219): Springer.
Ngram Viewer, G. (2021). Ngram Viewer. Retrieved from https://books.google.com/ngrams
Pattersons. (2022). PATTERSON ASSOCIATES LTD. Retrieved from https://pattersons.com/
Pazzani, M. J. (1991). Influence of prior knowledge on concept acquisition: Experimental and computational results. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*(3), 416.
Perlin, K. (1985). An image synthesizer. *ACM Siggraph Computer Graphics, 19*(3), 287-296.
Pharr, M., Jakob, W., & Humphreys, G. (2016). *Physically based rendering: From theory to implementation*: Morgan Kaufmann.
Pittman, K. (1992). A laboratory for the visualization of virtual environments. *Landscape and urban planning, 21*(4), 327-331.
Portman, M. E., Natapov, A., & Fisher-Gewirtzman, D. (2015). To go where no man has gone before: Virtual reality in architecture, landscape architecture and environmental planning. *Computers, Environment and Urban Systems, 54*, 376-384.
Quixel. (2022). Retrieved from https://quixel.com/megascans/
Remondino, F. (2011). Heritage recording and 3D modeling with photogrammetry and 3D scanning. *Remote sensing, 3*(6), 1104-1138.
Remondino, F., & El-Hakim, S. (2006). Image-based 3D modelling: a review. *The photogrammetric record, 21*(115), 269-291.
Render, E. (2020). How Architects can use Instagram to Their Advantage. Retrieved from https://easyrender.com/a/how-architects-can-use-instagram-to-their-advantage
Rhodes, A., & Wilson, C. M. (2018). False advertising. *The RAND Journal of Economics, 49*(2), 348-369.
RTA. (2022). RTA Studio. Retrieved from https://rtastudio.co.nz/
Sampson, H. (2021). *A history of advertising*: Prabhat Prakashan.
Samsung. (2022). Gear VR with Controller. Retrieved from https://www.samsung.com/nz/wearables/gear-vr-controller-r325/SM-R325NZVCXNZ/
Santibañez, D. (2019). Alexis Christodoulou: "Anyone Who is Qualified as an Architect Inspires Me". Retrieved from https://www.archdaily.com/920836/alexis-christodoulou-anyone-who-is-qualified-as-an-architect-inspires-me
Sapirstein, P. (2018). A high-precision photogrammetric recording system for small artifacts. *Journal of Cultural Heritage, 31*, 33-45.
Seaborne, M. (1995). *Photographers' London, 1839-1994*: Conran Octopus.
Seif El-Nasr, M. (2005). Intelligent lighting for game environments.
Seymour, M. (2013). Game environments – Part A: rendering Remember Me. Retrieved from https://www.fxguide.com/fxfeatured/game-environments-parta-remember-me-rendering/
Shashi, M., & Jain, K. (2007). Use of photogrammetry in 3D modeling and visualization of buildings. *ARPN Journal of Engineering and Applied Sciences, 2*(2), 37-40.
Statista. (2022). Number of Mobile Devices Worldwide. Retrieved from https://www.statista.com/statistics/245501/multiple-mobile-device-ownership-worldwide/
Steam. (2022). Steam Hardware Survey: October 2022. Retrieved from https://store.steampowered.com/hwsurvey/Steam-Hardware-Software-Survey-Welcome-to-Steam
Struth, M. S. (2014). *Clean (ing) coal: The enthymematic qualities of images, editing and color in clean coal advertising*: Wake Forest University.
Trends, G. (2021). Trends. Retrieved from https://trends.google.com/
Weber, E. A. (1980). Vision, composition and photography. In *Vision, Composition and Photography*: de Gruyter.
WetaFX. (2022). wetaFX. Retrieved from https://www.wetafx.co.nz/
Wilson, J. (2022). Physically-Based Rendering, And You Can Too!

facebooktwitter. Retrieved from https://marmoset.co/posts/physically-based-rendering-and-you-can-too/

Wilson, M. (2019). Welcome to post-reality: Apple will now fake your eye contact in FaceTime. Retrieved from https://www.fastcompany.com/90372724/welcome-to-post-reality-apple-will-now-fake-your-eye-contact-in-facetime

Zhang, Y. (2022). Algorithmic photography: a case study of the Huawei Moon Mode controversy. *Media, Culture & Society*, 01634437211064964.