Aalborg University Copenhagen

Semester: LiD 10

Title: Explore methods for digital asset reconstruction and their application within lighting design

Project Period: Autumn – Spring 2017

Semester Theme: Master Thesis

Aalborg University Copenhagen

Frederikskaj 12,

DK-2450 Copenhagen SV

Semester Coordinator:

Secretary:

Supervisor(s): Georgios Triantafyllidis

Project group no.:

Members: Iulian Drug

Abstract:

The following paper explores the topic of image capturing and digital reconstruction method known as photogrammetry. More specifically it focuses on the application of this method in the field of lighting design and aims to provide an in-depth descriptive account of different techniques and approaches of doing photogrammetry, and challenges that arise from them. The paper consists of multiple use cases, which help illustrate different aspects of the photogrammetry process, covering areas such as: image capturing strategies, de-lighting techniques, and digital workflows. It later transition into a discussion of automation and the advantages it can offer for the development of this particular field, and proposes a prototype design for an automated photogrammetry platform.

Table of Contents

Introduction

Lighting design is a field of work that nowadays bounds together art and engineering with the purpose of creating functional and aesthetically pleasing light schemes for interior and outdoor use. Before any luminaires are placed in the environment most lighting designers use different simulation software to calculate the energy use and visualize the light distribution. In the case of a new build, the structure, materials and light fixtures can be adjusted individually in the design development stage in order to achieve the desired result. However, in the case where the structure or the focus object already exists, accurate 3D models of the scene can be digitally reconstructed and the physical properties of the materials should be simulated.

The reasoning for this project started during the first weeks of studies in Lighting Design at Aalborg University in Copenhagen when we were introduced to a kick-off project concerning lighting design in a museum. Multiple lighting fixtures were accessible in order to plan a small lighting scheme, but the available time frame for testing in the space was very limited. For this type of context, where testing time is limited, the availability of a 3D model is important as different lighting schemes could be designed and simulated in software like Dialux.

Several months later, I had a talk with a curator from the Ny Carlsberg Glyptotek about a situation they were experiencing at that moment. A room from the museum had to be renovated and the lighting fixtures needed to be mounted before the statues and the artwork was placed inside. This situation gave the opportunity to ask myself the question of how would a lighting designer approach the task of creating a lighting scheme that captures the essence of an artwork, without having physical access to it?

In order to enhance each artwork with light and shadows, the lighting designer would benefit from having access to the 3D model of each sculpture and not only the room, as any simulation of a lighting scheme focused at organic shapes is dependent on having the intended target object in the scene. Thus, an iterative design process can begin for lighting schemes that emphasize the dynamism and shape of each intended element.

Furthermore, the advent of virtual reality and interactive design processes, can empower lighting designers to accurately visualize how lighting is affecting the materials used in the environment (Illumni.co, 2017). By using game engines as a design tool, the design process can become both interactive and iterative, as these tools are intended to permit rapid changes in the overall design scheme of an environment. For these reasons, I want to understand and explore different methods for digital asset reconstruction and their application within lighting design.

Currently there are a multitude of techniques for 3D data acquisition (inc., 2017), most of which employ a different type of remote sensing devices (sonar, radar, lidar, structured light, etc.) and the use of this data is diversified within each industry. Most of the data captured with these techniques comes in a form of point clouds of 3D coordinates which can be later computed in a 3D model. Depending on the sensor, for example lidar, color information of the target can also be recorded and further used. Though the results may prove to be very accurate, most of these techniques require a fairly large budget and specialized hardware and software used by trained professionals.

A secondary technology for 3D data acquisition and object reconstruction has risen in adoption in the past decade, called photogrammetry (GIS Lounge, 2017). It has its roots in aerial reconstruction, but it has been used with success even at the macro level. Photogrammetry involves estimating the 3D coordinates of points on an object employing measurements made from several photographic images taken from different positions, based on the principle of stereoscopy. Common points are identified on each image. A line of sight, or a ray, can be constructed from the camera location to the point on the object. It is the intersection of these rays (triangulation) that determines the three-dimensional location of the point.

Besides geometric reconstruction of the object using fairly low-cost hardware like DSLR cameras, the biggest advantage for using photogrammetry is the capacity to capture the color information or the texture of the object. Since lighting designers have the ability to use lighting fixtures with different correlated color temperatures, visualizing how that choice is affecting the material can prove beneficial for the overall design.

In the following paper, I aim to explore various methods for acquiring 3D models using photogrammetry and the steps required to use these assets for lighting visualization purposes in various real-time time engines like Unity (Unity, 2017) or Unreal (Unrealengine.com, 2017). The functional and economic rationality will also be taken into consideration, as these aspects affect the usability of any workflow.

Chapter 1 provides an overall in-depth description for the photogrammetry process in terms of image quality and capturing strategy.

Chapter 2 is focused mainly at the digital pipeline, where several software packages are used to transform the images into a 3D model.

Chapter 3 details the use of two different techniques for enhancing the usability of the texture maps by removing any residual light or specularity from the texture

Chapter 4 proposes a robotic system that aims to help with streamlining the capturing process and offers a solution for creating specular maps

Chapter 1. Capturing images for photogrammetry

The first step is to analyze how photogrammetry works and what are the main requirements for a successful digital reconstruction. I started by reading thoroughly the manuals for the major photogrammetry software providers since the principles of photogrammetry capture are universal and not bound to any specific software. Even though the software providers list is quite exhaustive (En.wikipedia.org, 2017), I chose to narrow the focus at 4 companies, representative in their market. The choice is based on product features, presentation material, availability of demonstrative models and user feedback from forums. The following short list details the software name and manuals for citation purposes:

- Reality Capture (Capturingreality.com, 2017)
- Agisoft Photoscan (Agisoft.com, 2017)
- Autodesk Remake (Autodesk Remake - Getting Started Guide, 2017)
- Bentley ContextCapture (Bentley Communities, 2017)

Compiling the similarities between all the recommendations, a set of key starting points was listed for two topics, image quality and capturing strategy.

1.1 Image quality

- Use the highest image resolution possible
- ISO should be set to the lowest value; high ISO values will introduce noise to images.
- Aperture value should be high enough for sufficient focal depth: it is important to capture sharp, not blurred photos.
- Shutter speed should not be too slow, otherwise blur can occur due to slight movements.
- Fixed lenses are preferred, especially 50mm. Avoid ultra-wide angle and fisheye lenses.
- Diffuse lighting is required to achieve quality results, yet flashes should be avoided

1.1.1 Photography Primer

When the camera captures an image, the exposure determines how light or dark that image will appear. Three camera settings control this: aperture, ISO and shutter speed. Photographers are calling this the ‘exposure triangle’ (Figure 1), because each setting influences different image properties (McHugh, 2016). The aperture affects depth of field, shutter speed affects the amount of motion blur, and ISO affects overall image noise.

The aperture determines the light gathering power of the lens, or more intuitively, it controls the size of the opening through which light can pass to the camera sensor. It is measured in f/stop, which is a ratio of the focal length of the lens to the diameter of the aperture (Verhoeven, 2016). Every f/stop doubles the amount of light passing to the sensor, and though counterintuitively, the smaller the f/stop number the larger the aperture.

The aperture affects the range over which objects appear sharp both in front and behind the point of focus, creating what is known as the “depth of field”. Increasing the depth of field is essential to maximize the amount of usable data for reconstruction, as any spot that is not in clear focus will either be disregarded by the reconstruction algorithm or contribute to the overall noise on the geometry. There is a limit for decreasing the aperture, because of a phenomenon called diffraction, where the light rays passing through a small aperture will start to interfere with one another creating an overall softness and blurring the resulting photograph.

For example, in this test (Table 1) several images were taken by a Nikon D800 with a Micro Nikkor 60mm f/2.8G ED lens. Reality Capture was used to analyze the photos and provide a list of detected features per image. Macro lenses are generally considered to be very sharp lenses even at extremely low apertures, but starting from f/29 diffraction affects the overall quality of the image.

Aperture	f/5.6	f/7.1	f/9	f/11	f/16	f/22	f/29	f/36
Reconstruction Features	25131	33691	42793	54703	72807	99275	100161	78633

Table 1 Diffraction limit test on Micro Nikkor 60mm

In a DSLR, the shutter is a curtain that blocks all incoming light and stays closed until the shutter button is pressed. Then it opens, the sensor is fully exposed to the incoming amount of light and then it is immediately closed again, blocking the light. The duration that the sensor is exposed to light is called the shutter speed. Generally, it is desired to have a fast shutter speed to freeze the action, because slow shutter speed can create motion blur on the whole image.

The ISO speed is figure that determines how sensitive the sensor is to incoming light. The lower the ISO setting, the less noise will be found in the image and the higher the image quality (Verhoeven, 2016). ISO is also very dependent on the camera sensor manufacturer, because one of the biggest difference in between various camera manufacturers is the so-called SNR – sensor to noise ratio.

Higher end cameras, like Nikon D810 and Canon EOS 5D Mark IV have sensors with higher sensor sensitivity in low light, allowing the photographer to increase ISO to higher values, like ISO 800 or ISO 1600, without introducing a comparable amount of image noise. (Dxomark.com, 2017)

I mentioned earlier about the exposure triangle where one needs to consider camera settings to balance each other’s strengths and weaknesses, because changing one setting will affect the other two and everything depends on the amount of light present at the moment of capture. During the last year several dozen tests were performed to try to find a balanced approach for image quality and quantity when concerning photogrammetry. To document the procedure for all these tests is not the scope of this paper, but I would like to include findings based on my personal experience working with this topic.

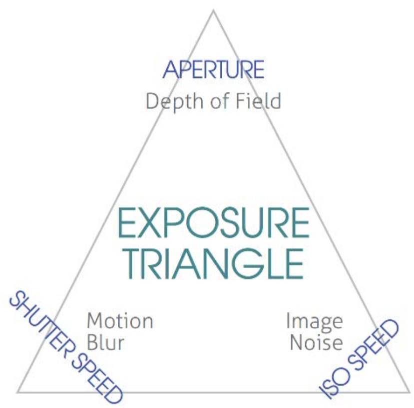

Figure 1 Exposure Triangle

In my opinion, for handheld photography, everything starts from the shutter speed. This must be set to diminish any possible motion blur, then the ISO and aperture must be set in tandem to each other, raising or lowering one than the other to decrease the risk of having either too much noise or too little depth of field. Working in a controlled environment, the wheel turns around and ISO and aperture can be set to optimum values dependent on the camera and lens, for example ISO 100 and f/8 – f /11. The shutter speed can be set to a lower amount with the use of a tripod or a fixed base and / or artificial light can be added for proper exposure.

Good resources on technical aspects of photography can be considered (Peterson, 2004) and (Salvaggio, Stroebel and Zakia, 2013). Also, a great online resource is Cambridge in Color website that has several tools for calculating camera settings according to different situations, some of which also support my findings. (Cambridgeincolour.com, 2017)

1.1.2 Resolution

Commonly, the pixel count or the total number of pixels is referred to as the image resolution, but this does not present a clear statement of the camera's ability to distinguish fine detail. For example, a camera with 36-megapixel sensor can capture a photo of a coin on the table, but pointed outside the window and it can capture an image of the whole neighboring building. The 2 images, although both counting 36-megapixels feature a very different spatial resolution or pixel to meter ratio.

The relationship between physical detail and the camera's ability to capture that detail can be calculated by the spatial resolution equation.

$$R = \frac{Ls\ x\ D}{f\ x\ L}$$

R is the spatial ground resolution in meter/px
Ls is the width of the sensor in mm
D is the distance between the camera and the object in meters
f is the focal length of the lens
L is the width of the image in pixels

A simple test (Table 2) was devised to check for the relation between sensor size and spatial resolution. Two different cameras were used, with the same settings and lenses chosen for comparable spatial resolutions. In total, 40 photos were taken with each camera from the same physical point, approximately 50 cm from the object. Artificial lighting was provided in the form of studio flash light diffused with a softbox. Reality Capture was used to analyze the photos and provide a list of detected features per image.

Camera	Lens	ISO	Aperture	Shutter speed	Spatial Resolution	Minimum Features	Maximum Features	Total no. Features
Nikon D800	60mm f/2.8	100	f/8	1/200	24.6 px / mm	1718	14693	161318
Nikon D3300	50mm f/1.8	100	f/8	1/200	25.5 px / mm	1642	13918	153214

Table 2 Spatial resolution test

This test should not be seen as a de facto camera performance test, but more as guide to exemplify that the sensor size is not the defining factor when it comes to spatial resolution and the ability of the camera to resolve detail. These two particular cameras are at the different price ends of the Nikon DSLR collection, but can deliver comparable results when the right camera components and optics are working together properly to deliver a high-quality image.

Professional camera performance tests are done by several institutions around the world, with Dxomark being considered one of the best for consumer level (Dxomark.com, 2017). They specialize in benchmarking cameras from different manufacturers and providing camera – lens specifications that are also used in their software Dxo Optics Pro. This can be a good starting point to assess camera + lens combinations that yield good results.

1.1.3 Lenses

Lenses are classified by their Field of View into three main categories: wide-angle, standard and long focus or telephoto lenses. Their classification takes into account the human vision system, where a static human eye can have a sharp FoV of approximately 50^{o} (Verhoeven, 2016). The photographic term used to express the FoV is focal length, where lower numbers mean wide angle lenses(24mm), and bigger numbers telephoto lenses(200mm).

Wide angle lenses have a typical FoV wider than 60^{o}, with extreme wide-angle lenses going all the way to 120^{o} and fisheye lenses to 180^{o}. Standard lenses cover a FoV between 40^{o} and 60^{o} and telephoto lenses from 35^{o} all the way to 1^{o}. What is important to note is that dependent on the size of the object, the distance to that object and the required spatial resolution the lens with the appropriate FoV shall be chosen.

As an example, for a required spatial resolution of 10px/mm, I computed this table (Table 3) to express better the relation between the focal length / FoV and distance to the object.

Distance (m)	0.3	0.5	1	2.7	6	10
Lens (mm)	16	24	50	135	300	500
FoV (degrees)	110	84	53	18	8	5

Table 3 Varying lens choice based on distance and spatial resolution

Besides the desired spatial resolution and the distance to the object there are two more important factors to consider when choosing a lens. When used handheld, telephoto lenses require a high shutter speed to avoid motion blur. The general rule-of-thumb is that the shutter speed should be equal or faster than the reciprocal of the focal length in use. As an example: a 300mm lens necessitates a shutter speed of at least 1/320s (Verhoeven, 2016).

The second factor that needs to be taken into consideration is the lens distortion. The general advice is that only rectilinear lenses should be used for photogrammetry. Rectilinear lenses render straight lines as straight lines, opposed to fisheye lenses that display them as curves.

Even though every tested photogrammetry software has a database of lens distortion factors that is used to undistort the image, the resulting reconstruction still suffers from it. For example, a straight part of a wall was photographed with a Nikon D800 with a 16mm fisheye lens and 50mm lens. The resulting reconstruction (Figure 2) clearly displays the straight wall as being curved, even though several steps have been taken to undistort the image.

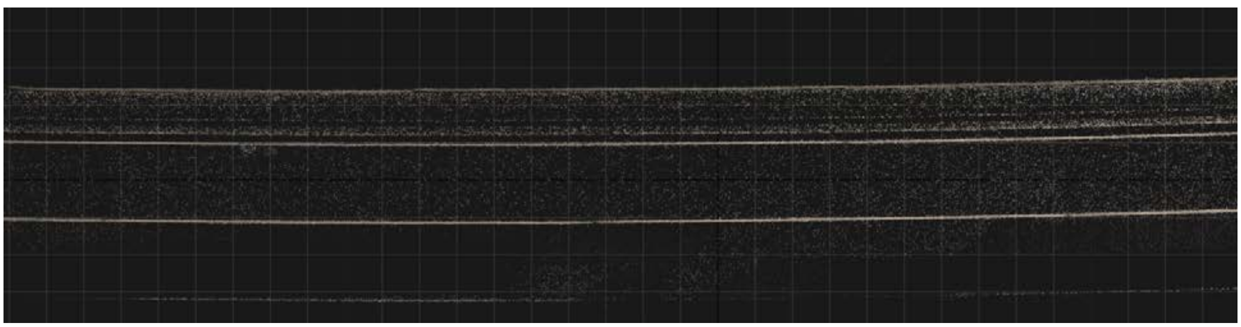

Nikon D800 with a Nikon 16mm f/2.8 fisheye lens

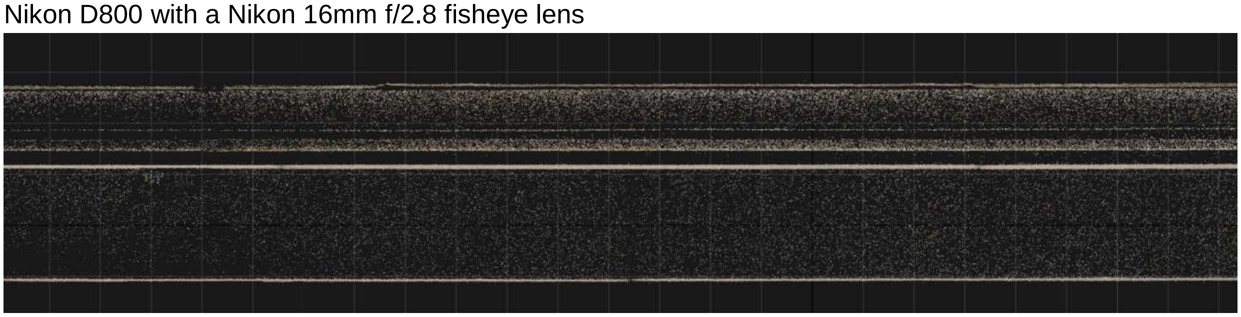

Nikon D800 with a Sigma 50mm f/1.4 Art rectilinear lens

Figure 2 Distortion test

1.1.4 Lighting

All photogrammetry software development companies advise that diffuse lighting is required to achieve quality results. A lighting source can be considered diffuse if it scatters or does not have a clear direction. Furthermore, it does not create hard shadows, but seems to wrap around the object. The skylight, without clear sunlight, is the best example for diffused light, and also the most recommended type of light when it comes to photographs for photogrammetry.

Since it is not possible to always have diffused daylight as our main light source, there are several options to recreate the same effect with artificial light, most of which employ some type of diffuser, usually frosted glass or fabric. The size and distance of the light source can be also taken into consideration when discussing about light diffusion because the same effect can be attained by either having multiple smaller light sources or larger one but further away. This is due to the inverse square law or the decay of light over distance with a common use in this context: double the distance between the object and the light source and it illuminates a surface area four times larger than before.

Photography flashes or strobes are recommended to be avoided, but there is no clear stated reason for this. In order to test what exactly is the problem with using flash/strobe lights, several tests were conducted on a color profiling tool, the X-rite ColorChecker (Xritephoto.com, 2017). To analyze the mean color deviation between different shots I used 3DLUTCreator (Sharonov, 2016) because it allows the calculation of different image metrics and also supports a wide variety of color targets.

I assumed that this recommendation has to do with the fact that flashes are electric arc lamps filled with some type of noble gas, like xenon. They are designed to produce intense, full-spectrum white light for very short duration. The issue with this type of lamps is their energy and color temperature stability. Another way to look at this is the probability that at each trigger it discharges exactly the same amount of lumen with the same spectral distribution.

I decided to test the Godox AD360II-N WITSRO TTL and the Profoto D2 1000 AirTTL. These two flashes were chosen due to availability and large difference in between price points. In order to remove as much bias from the data as possible, both flashes were set to manual control and the power output was measured using the Sekonic L-478D exposure meter. The flashes were then adjusted for an even exposure value.

Since luminance and color consistency is tested between two very different flash units, I concluded that the test should be done in comparison to how the natural light behaves. The ColorChecker was placed outside in two different places: in unobstructed light from the sky and also in the shadow. In total, for each scenario, ten photos were taken with an interval of 1 sec between shots. Afterwards the RAW files were processed linearly and the values from 3DLUTCreator were computed in an excel sheet.

Lighting conditions	Average Deviation	Maximum Deviation	Reconstructed Features percentage deviation
Daylight	0.06	0.12	0.267%
Shadow	0.07	0.26	0.415%
Godox	0.7	0.9	2.532%
Profoto	0.15	0.36	0.441%

Table 4 Color consistency test

The table (Table 4) presents only the resulted average and maximum deviations from the measured values on the ColorChecker chart. The results enforce the assumption that flash photography is not desired if it presents color temperature instability. Furthermore, Reality Capture was used to analyze the number of detected features between each image. The resulting percentage deviation suggest that a high-quality flash module, like the Profoto D2, has a good luminance and color consistency closer to daylight values and better suited for photogrammetry than the Godox flash module.

1.2 Capturing strategy

Every object is unique in its own way and may require a different approach for a successful reconstruction. Nonetheless, a general capturing strategy is a good starting point. An object's physical properties determine what is important to consider before actually starting taking photographs. Considerations should be made for the overall size of the object, the available distance, the material of the object and the required precision. The following list, based on reviews on the manuals and personal experience, can be considered a general guideline for photographing most objects:

- Start with pictures of the whole object, then move closer and focus on areas that require more detail
- ensure a minimum of 60% overlap between images
- photograph the same part of an object from at least three different points of view
- Limit the angle between two consecutive photos to a maximum of 15°
- Each photo should effectively use the maximum frame size
- Avoid specular or transparent objects like glossy plastic or glass

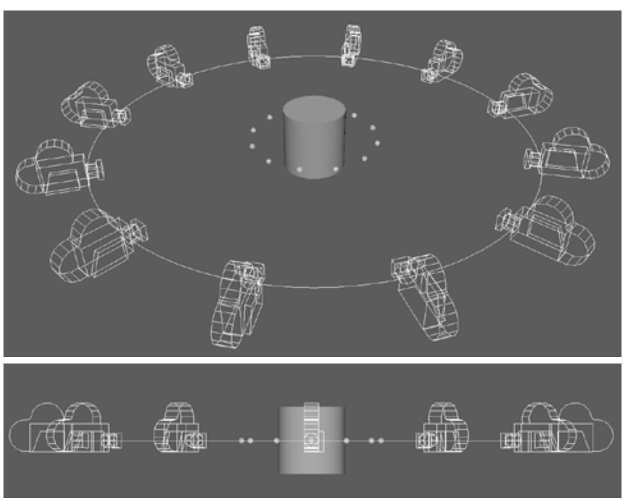

Figure 3 First image set around the whole object

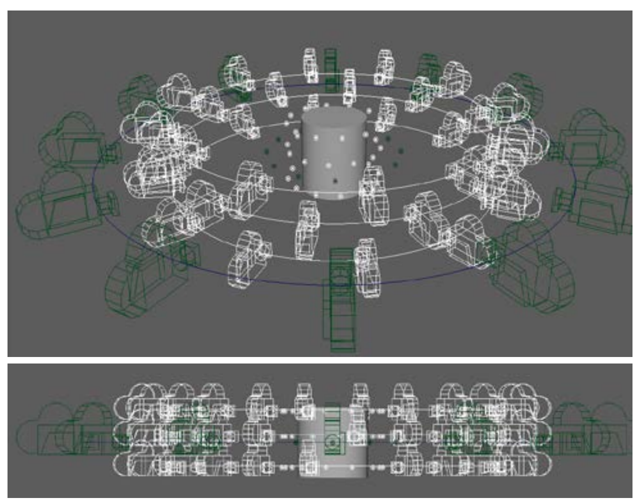

Figure 4 Close-up photographs paralel to the surface

Analyzing the recommendations, we should start with pictures of the whole object, then move closer and focus on areas that require more detail for the final 3D model. The first set of pictures (Figure 3) can be considered the outer layer that helps to recreate the rough shape of the object and also to link with the next, more detailed set of pictures (Figure 4). It depends on the object, but most of the time several layers of pictures are needed to increase the chance for reconstruction. When it comes to straight edges, for example the cap of the cylinder, at least one set of pictures must be focused on the edge itself (Figure 5), with as much depth of field as possible. This way the pictures from the two sides (Figure 6) can be linked together by this middle set.

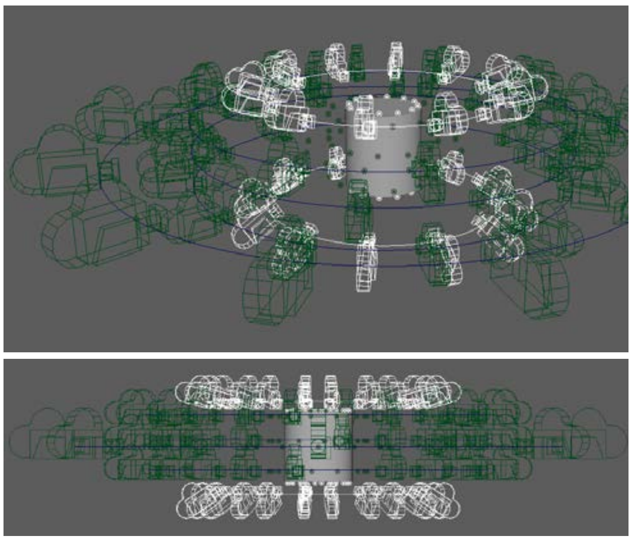

Figure 5 Focus on the edge

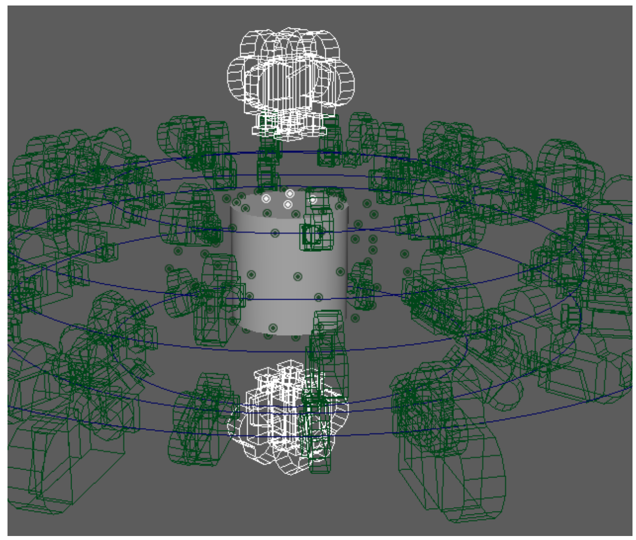

Figure 6 Top and bottom angles

1.2.1 User case: The Thinker statue by Rodins

The Thinker statue by Rodin, placed in Ny Carlsberg Glyptotek garden, is approximately 2 meters tall, mounted on a stone pedestal also approximately 1.5 m tall. It is closely surrounded by flowers and small bushes, making the close-up capture possible at a minimum distance of 1 meter from the statue and mainly from two sides, front and back.

The following list details the equipment used for capturing the statue:

Camera: Nikon D800 full frame DSLR
Lenses: Sigma 50mm F1.4 Art, Nikon 300mm f/4 AF
Monopod: Benro MMA38C Mach3 9X Carbon Fiber Monopod
Extra: X-Rite ColorChecker Passport

The capturing process started by hand holding the camera with the Sigma 50mm lens and making several rings of capture around the object, moving after each shot and changing the focus point constantly between shots to make sure that all the visible parts are photographed from several points of view and different elevations (Figure 7). The ISO was set to 100, aperture to f/11 and shutter speed to 1/200.

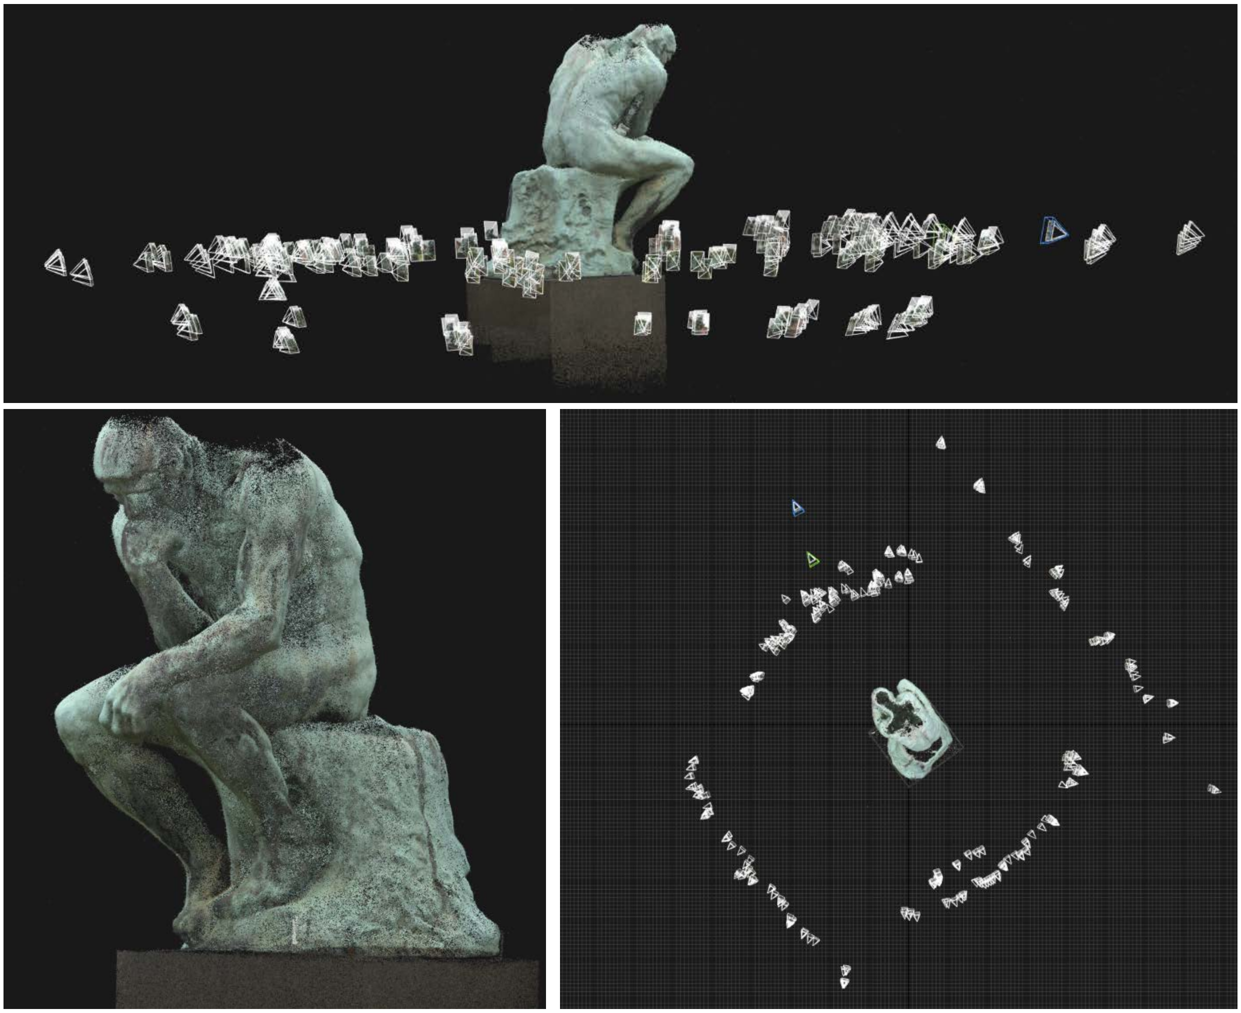

Figure 7 First set of photographs with Sigma 50mm

After this initial set of 330 pictures, it was clear that close-up pictures need to be taken from steeper angles especially for the areas under the arms, palms and chin and also an entire set for the top. The camera was mounted on the monopod, the shutter speed was increased to 1/250 to decrease the chance for motion blur, the ISO kept at 100 and aperture changed to f/9.

Since I did not have an external monitor with a real-time view from the camera, I set up an interval shooting for every 3 seconds. I set up a batch of 50 images at a time and started moving the monopod closer to the interest areas while at the same time approximating where the camera was automatically focusing, waiting for the shutter trigger and moving after each shot. At the end of each set I went through the pictures, making mental notes where the coverage was not enough, adjusted the inclination angle of the camera on the monopod and started again the process.

Figure 8 Second set of images using a monopod

In total, I took 320 more images with the monopod. The second set with close -up images brought in more details in the occluded areas and on the top of the statue, but still required some more photographs from the top (Figure 8). Unable to reach that high, I changed the lens to a 300mm telephoto, adjusted the camera settings to this lens with a shutter speed of 1/320, aperture of f/5.6 and kept the ISO at 100. I positioned myself several meters away on top of a small wall and took some more pictures for the top side (Figure 9). I walked around one last time and took in total 100 more shots to make sure I had enough pictures and a good coverage on the overhang areas.

Figure 9 Last set of photographs including the Nikon 300mm lens

Reality Capture has a very good tool called Inspect (Figure 10) that can help visualize using the blue lines the connections between all reconstructed features of all images. This can also be considered a debugging tool as well since it also shows the area where there is a low coverage, like the top part of this statue.

In practice, it is always harder to maintain the necessary overlap and it easy to get distracted and forget elements during the shoot. The most important part is to have full coverage of the object, preferably by several pictures from different points of view with a high overlap.

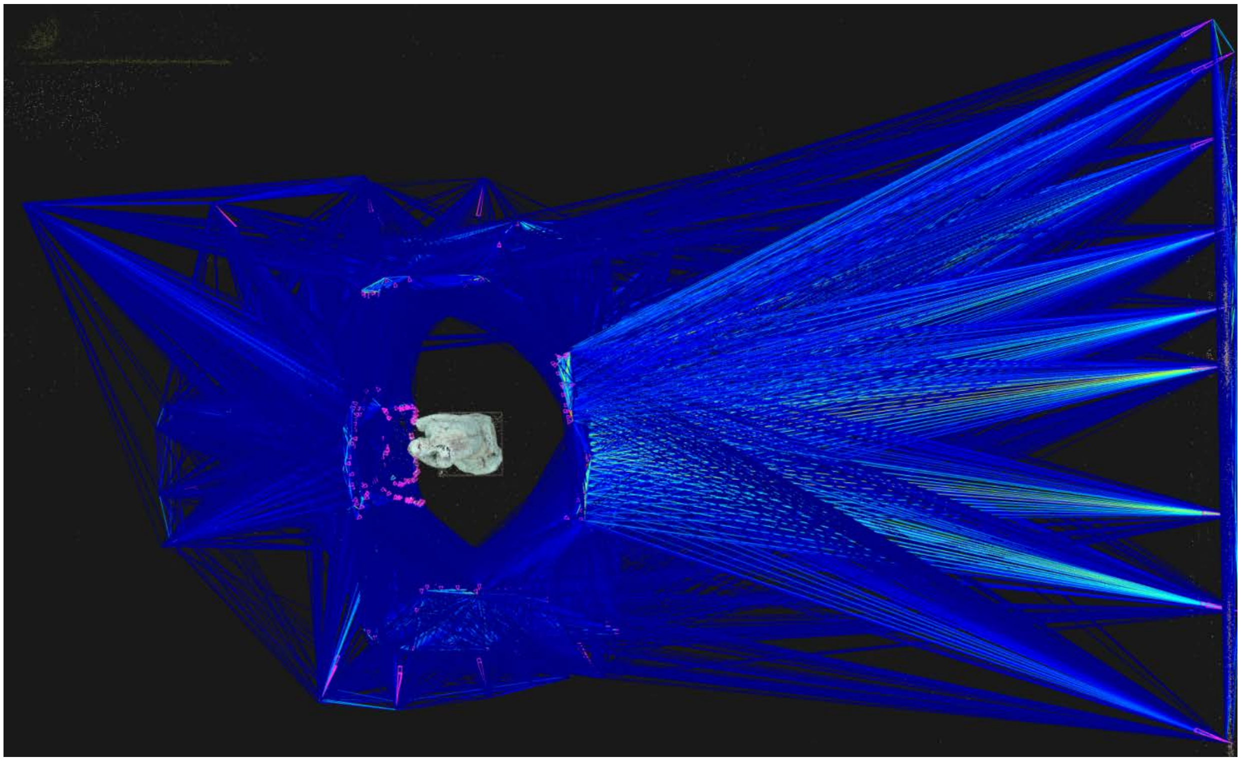

Figure 10 Inspect tool in Reality Capture

An important aspect when conducting a capturing session is the correct color consistency throughout the picture set. The X-rite ColorChecker can be used as a scene reference for white balance and color correction, but also for scale. An inherent disadvantage with photogrammetry is the fact the scene scale is not known, or it cannot be deduced automatically from the images. However, there is an option to set the scale of an object, if a prior distance measurement has been taken in the reconstructed scene. An example is a physical ruler or the ColorChecker next to the object that is to be reconstructed.

Chapter 2 – Digital pipeline workflow overview

When dealing with real-time and/or high-quality renderings, there are steps in the pipeline where small differences or errors can occur, such as if the captured light accurately represents the light levels in the scene, or the renderer's capability of simulating a real environmental light. As such, in the following section, image processing terminology and methods related to photogrammetry will be discussed.

2.1 – Image processing primer

Any display that we use to view images does do not linearly convert voltage into light intensities, but the feedback is represented by a curve which applies an exponent value or a gamma of 2.2

to the image and it affects the intensity of each pixel. The adjustment of the pixel values is called gamma correction. A linear workflow defines the manipulation of gamma correction, to ensure an improved and accurate image display. (NVIDIA Developer, 2016)

Firstly, light transport is linear as two light sources are additive and not multiplicative and do not interfere with each other in any other way. The problem when working with renderers is the input data, mainly the input images. If the captured image is nonlinear and it is rendered linearly, then the result might look natural on a display, because it may cancel out the effect after the image is processed. However, in that case, the RGB color values of pixels are not properly calculated and there will still be subtle differences in the pixels that represent an unnatural behavior to changes in brightness and colors.

In order to avoid obtaining artifacts, gamma correction needs to be taken into account, however, the problem first arises when the pipeline includes nonlinear input textures. Most of the images captured are already corrected and the color values need to be changed back into a linear space. An ideal solution to obtain a better workflow is to use a different file format, that has the possibility of capturing an image and not correct it. There are three main data types for graphical storage and representation. Raster data, vector based and latent image data (Verhoeven, 2016). In this discussion, three file formats will be presented. RAW image, which is latent data and has the metadata that needs to be processed into raster formats, and TIFF and JPEG which are raster formats.

The RAW format is represented by the unprocessed raw data that it stores. It contains the analog information transformed into digital data, without being altered by the camera's software and it is the base form of photographic data. Because it contains many file formats, there is no standard format for all manufacturers, however, it can still be used for processing, as it contains the metadata information. There are two pipelines that can process this information. One is done by the cameras image processing engine and it determines the final output of the image. The other one is done via computer by the user. There is a numerous amount of software systems that can process a RAW file and they offer a wide range of options for modifying specific parameters in the file, such as, color correction, noise removal and sharpening.

The TIFF file format (or Tag Image File Format) stores the image in a body section and also creates tags, which store the data in the original order, such as, the color profiles or the depth of the image. Also, it supports 16-bit data information per color channel and it is the preferred format to compress to, from the RAW format. This being the case, the file size is large, as it stores all the information. Even though this file format is very close related to the RAW format, it is unsuited for capturing images with a camera as the file format will be larger and the scene data is not the original captured information, because it is altered by the camera's internal engine system.

The last file format in this discussion is the JPEG (or Joint Photographic Experts Group) format, which is not composed of one compressing algorithm, but a multitude of algorithms with different capabilities (Verhoeven, 2016). The main compressions used are, the lossy, which compresses the file as much as possible, ignoring a lot of information, but still containing an almost undistinguishable visual effect from the original input file, and the lossless compression algorithm which keeps most of the information, but it compresses to a smaller factor than the first one. The latter algorithm is used in most applications, as it has the possibility of compressing a file to a very small size compared to the original. Also, this compression can only store 8-bits of information per color channel and a camera's internal engine system applies a tonal curve to the captured image, cutting down on the original scene information, therefore, as in the case of the TIFF file format, it is unsuited for capturing images.

When talking about bit depth, the smallest values that a pixel can get are 0 and 1, respectively being black and white, meaning the pixel has 1 bit of information. By increasing the amount of bits of information, more possible shades appear for the pixel. An 8-bit image contains 256 shades of color per channel. The problem when dealing with this type of images is that, if an enhancement is performed, then there is chance for banding in the image, which creates posterization, and it doesn't make the transitions between gradients smooth enough. If a 16-bit image is used, which contains 65,536 shades per channel, then the transition is much smoother, as the gradients have more possible combinations. In HDRs, 32-bit images are used, with 4,294,967,298 shades per channel, in order to obtain the best transitions in tonality of shadows and highlights(D&M Imaging, 2017).

2.2 Physically Based Rendering

Physically Based Rendering or PBR, is mostly considered a methodology for authoring texture maps that can used in PBR rendering engines. Even though the research principles of Physically Based Rendering have been defined around 2004 (Pharr and Humphreys, 2004), most rendering engines started implementing this starting from 2014 (Sébastien Lagarde, 2015). Each rendering engine, aimed at stills like Chaos Group V-Ray or real-time like Unity or Unreal Engine, implement this model now with various differences, adopting both or metal/roughness and specular/glossiness workflow.

The main difference in between the two workflows is that the metal/roughness workflow necessitates the authoring of a metallic map, which acts as a mask to differentiate between metal and dielectric data found in the base color map while the specular/glossiness workflow uses both maps to define the material’s response to incident light. (allegorithmic, 2016). While both approaches have comparable results, for the purpose of this project I will use the specular/glossiness workflow, because it offers full control on authoring the maps without being bound to hardcoded values in the PBR shader.

2.2.1 Texture map authoring guidelines

Diffuse Albedo Map (sRGB)
The terminology for this map might be different depending on the rendering engine, (be it diffuse, albedo, or color map) the diffuse albedo map contains the base color of a material, and

most of the time this is called the albedo color. This map should be void of any directional light and ambient occlusion as the material will look incorrect in different lighting conditions.

Glossiness Map (grayscale linear)
The glossiness map depicts the microsurface irregularities of a material that cause the amount of light diffusion. Depending on the material type, rough surfaces will present wider and dimmer specular reflections while smooth surfaces have brighter specular reflections or highlights.

Specular Map (grayscale linear)
The specular map defines the reflectance values for both metals and dielectrics. Dielectric materials have all the color information in the diffuse albedo map and grayscale specular values in the specular map while metallic materials have the color information in the specular map.

2.3 Processing the images

Following up with Thinker statue use case, the next section will provide a detailed description for the process of converting the photographs into a 3D model. As it is important to maintain a linear workflow throughout the pipeline, RAW photographs from the camera must be first converted to uncompressed linearized images. DCraw is generally considered the best option since it does not apply any contrast curves to the final image like Adobe Camera Raw or DXO Optics Pro (Guillermoluijk.com, 2017).

To convert the RAW images, the command prompt needs to be run with the following command "dcraw -v -4 -T *.NEF" in the folder containing the images and the DCraw executable. The options mean the following: -v is to display textual information about the process, -4 is the option to convert to 16-bits linear with camera white balance and a gamma of 1 -T is the conversion to uncompressed TIFF format. The next step is to white balance the photos using the X-Rite ColorChecker also save the files for further processing.

2.3.1 Reality Capture

The reconstruction software is the most important part of the digital workflow since it is the software that can take as an input all the images of an object and output a 3D model. I have chosen Reality Capture for this project because of the overall efficiency compared to Agisoft Photoscan and Autodesk Remake. During the comparison tests, Reality Capture was in between 2 – 10 times faster than both other solutions and the reconstruction presented more details in areas that with shadows or highlights.

The first step is to load the images and align all of them. If the alignment process fails and created several components with less than 50% of the total images, a way to fix that is by loading sets of images at a time. For example, all the images from the outside layer can be aligned first and after that step by step the rest of the images. If this still does not work, control points can be added manually to help the alignment procedure. If the procedure fails multiple times, most likely the object needs more photos. After the alignment is done and the resulted tie points look even overall, the next step is to select a reconstruction bounding box around the interest area. A failure to do so will force the software to try and reconstruct the whole scene (Figure 11) that contains the general environment. In most cases this will also cause the software to crash and the whole process needs to be started again.

Figure 11 Reconstruction region

Figure 12 Scale reference

The last step before reconstructing the object is to set the scale. I placed the X-Rite ColorChecker next to the statue in the scene and used the scale constraints in Reality Capture to specify the measured distance between those specific points (Figure 12).

Next, the reconstruction quality should be chosen. The difference in between the preview, normal and high-quality reconstruction is the amount of down-sampling of the images. If the highest precision is needed, high quality with no down-sampling should be chosen, however processing time will increase dramatically. The total project time can take in between 2-3 hours to several days depending on the number of pictures, resolution and available hardware.

Figure 13 Point cloud dataset

After the computer has finished processing, the reconstructed mesh is previewed as a point cloud (Figure 13) generally split in multiple parts of approximately one million points. In this case the final model is split in 54 parts totaling 38 million vertices. There are just a handful of software packages that can work with such a large dataset (Geomagic.com, 2017), and in most situations, it is not a good idea to bring high detailed models in rendering software or game engines. Instead, low polygon models are created. The low polygon model is a decimated or simplified version of a 3D model that contains the data for the overall shape, but lack any of the high frequency details of the high-polygon model. To maintain the same detail impression as the high-polygon model, special kind of textures, called normal maps, are created to add the surface details to the low polygon model and modulate how light will influence the surface (Docs.unity3d.com, 2017).

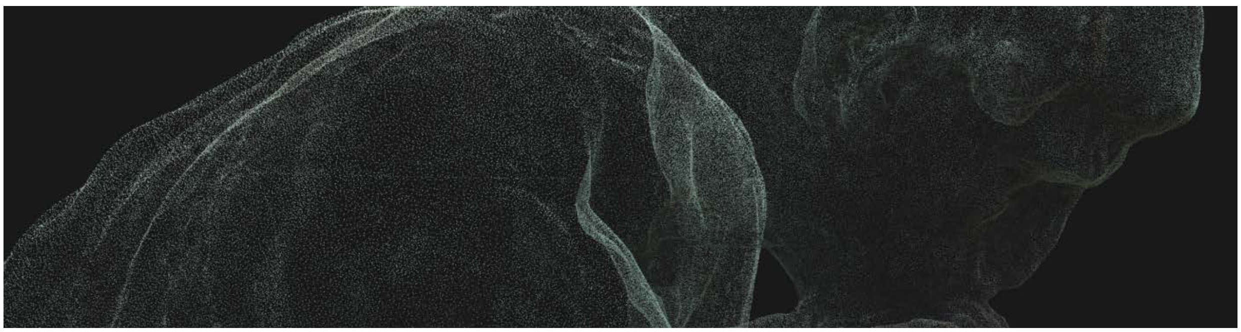

Figure 14 Decimated dataset

To create the low polygon model version of the statue, we use the simplify command directly in Reality Capture. This will take the overall shape of the model and reduce the polygon count to the desired value (Figure 14). Through trial and error, I noticed that models that are simplified under 2 million polygons start losing too much detail and displaying artifacts, and in this case, we need to use another software to further decimate the model.

2.3.2 Pixologic Zbrush

Zbrush (Pixologic ZBrush, 2017) is considered the industry standard for digital sculpting, and has been in active development for more than 15 years. It features a large set of tools that can help with our task of further decimating the mesh, while maintaining the outer shape. A secondary, free option, is represented by Instant Meshes created at the Disney Research Labs in Zurich. In my experience, both are well suited for the task of decimating the model, but I personally prefer working with Zbrush.

After exporting from Reality Capture the 2 million polygon model, I imported it to Zbrush and used the Decimation Master plugin to further process the model to 65.000 polygons which is also the maximum amount that Dialux can handle importing per 3D model (Figure 15).

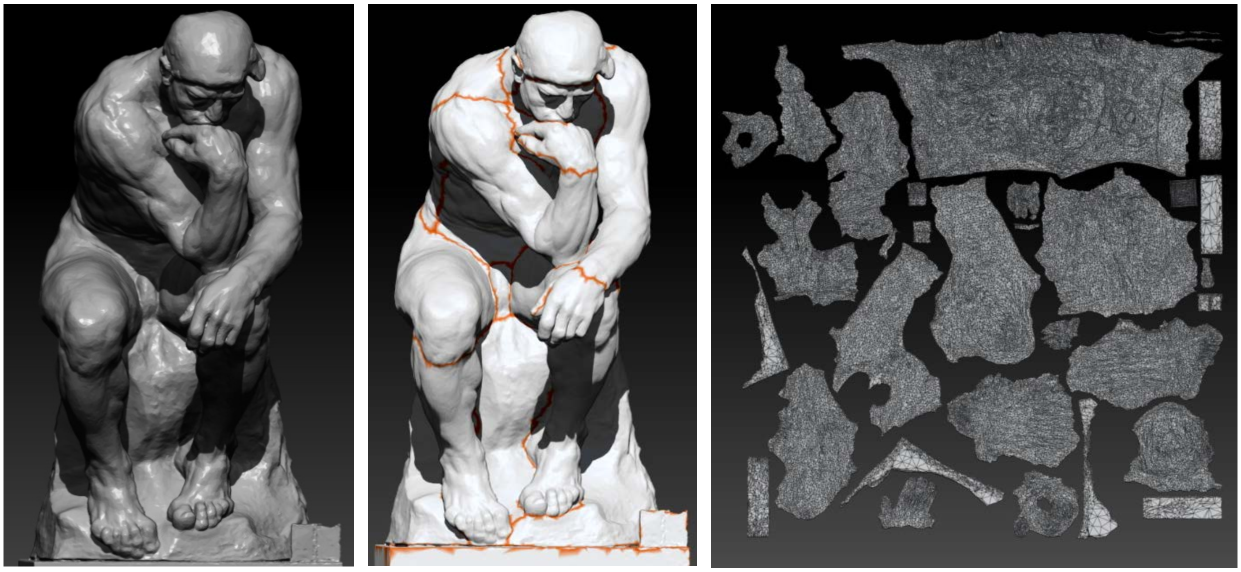

Figure 15 Decimated low-poly model with UV map

Before creating any textures, the UV map must be created. A UV map can be considered a 2D projection for a 3D model. If we consider for example the Earth, the Mercator projection world map is the UV map of the Earth. There are several ways, both manual and automatic to create UV maps, but for this example it has been created manually in Zbrush (Figure 15). In a nutshell, the procedure works by creating seam lines across the model and then flatten and unwrap the resulting patches.

The next step is to get the texture information from the model. This step is done in Reality Capture, where the low poly model containing the UV map is imported and textured with preferably the highest texture size that one needs. For example, on this project the exported texture is 8K texture, but depending on the hardware configuration, textures with a smaller size might be needed (Figure 16).

2.3.3 Xnormal

To bake the high detail information into a normal map we use the free software Xnormal (Xnormal.net, 2017). A good reason for using Xnormal is the fact that it supports very high polygon models, for example the maximum size I was able to load was almost 300 million polygons. It has a simple workflow where the low poly model containing the UV map is imported together with the high polygon model exported from RC. Then the software projects all the small details onto the low polygon version and then everything is saved or baked onto a 2D texture (Figure 16).

The resulted assets after doing this process are: the original reconstructed high-polygon model, a low polygon version of the model with a normal map and the original texture of the model.

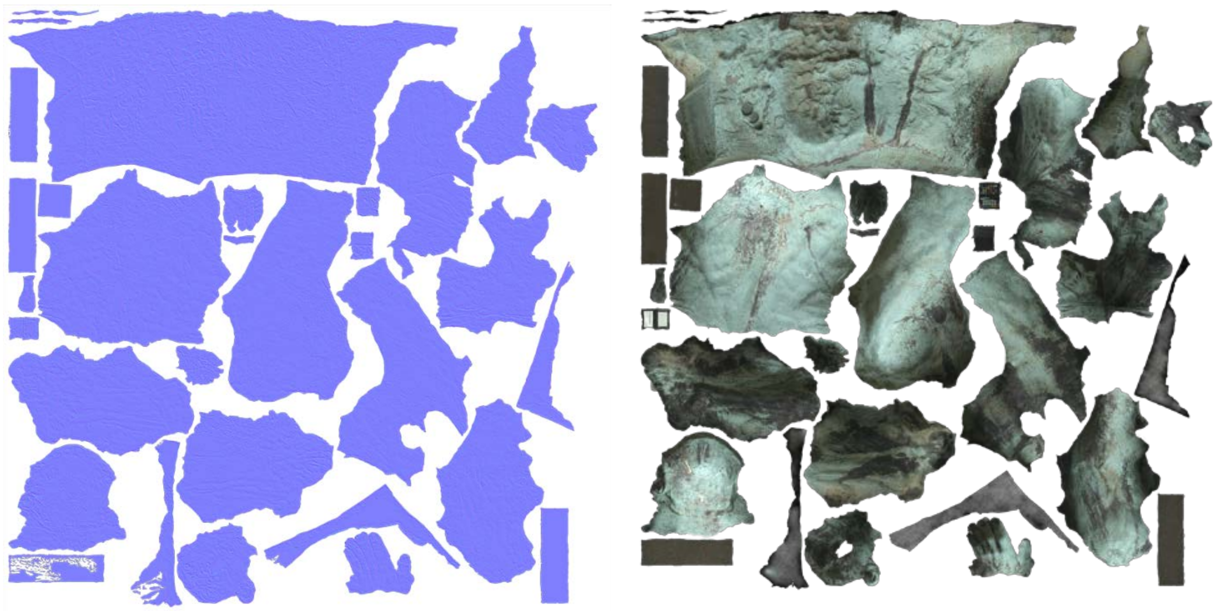

Figure 16 Normal map and texture exported from Reality Capture

However, during the capturing process, it was evident that the light in the scene will affect how the final texture will appear. In order to use the model in Unity and dynamically lit the object we need to create a diffuse albedo map that has all the lighting information removed from the texture.

Chapter 3. De-lighting

De-lighting is the common name for a process involving the removal of baked light information from texture maps and it has been in development at different research institutions and VFX companies for the last couple of years. One of the first research projects into removing light information from scanned object has been done by Paul Debevec and co. where a set of high dynamic range images were taken at the scene and used afterwards to perform an inverse global illumination using a custom renderer (Debevec, 2004).

The latest successful approach for a standalone software was done by the company Tandent (fxguide, 2017) with their software Lightbrush. However, the project has been discontinued now, and is no longer available. Several other approaches (Gdcvault.com, 2017) have been used to de-light a texture by using the ambient occlusion map and the object space normal map to remove the shadow information. The ambient occlusion map is calculated based on an approximation of how exposed each part of an object is to ambient light. For example, the inside of a tube is less exposed to ambient light than the outside and hence it is darker. Since the object space normal map displays the direction for each of the objects' normals, it can be used artistically to control the shadow levels.

Using any or both approaches to remove the lighting information from the diffuse albedo map can give good results in some cases where the reconstructed object is mostly planar (Gdcvault.com, 2017), but it does not take into account the lighting conditions in the scene at the time of the capture. While it is common to use synthetic representations of sky models in physically-based renderings, this does not represent the real-world lighting conditions (Kumaragurubaran, 2013). To accurately record daylight variations at the time when the object was captured, high dynamic range photographs of the environment can be used as the source of illumination for objects in a 3D scene within a process called IBL or image-based lighting.

Photographs from most consumer level cameras are considered low dynamic range images because the amount of information they can represent is limited. A dynamic range is the ratio between the largest and smallest values than can be recorded by a sensor at one time. A HDRI or a high dynamic range image is a combination of different photographs taken with different exposures. Since the luminance level in a scene can vary well beyond the capabilities of the sensor, the HDR photography technique has emerged as both an artistic and measurement technique to capture luminance values with higher dynamic range.

HDR images used for IBL generally need to be omnidirectional, sometimes called spherical or 360, as light originating from all directions contribute to varying degrees to the appearance of the real-world object. HDR images used for IBL contain information about the intensity, shape and color of all light sources, be it direct or indirect bounces from the surfaces in the environment. This method can be used to record the light information from the scene as a spherical luminance map and use that to balance out the lighting information in the diffuse albedo map just by using a photo editing software.

3.1 hdrscope

For de-lighting purposes, we need to capture the luminance map of the environment and use it to balance the shadow and highlights information from the diffuse map. Since we are using a linear workflow and we want to match the exposure, we will need to make sure that the HDR is photometrically calibrated.

During the studies, we were introduced to luminance mapping techniques for glare analysis that use Radiance software (Radiance-online.org, 2017) in the backend for generating calibrated false color luminance maps. Photometric calibration is done by either using luminance or illuminance measurements in the scene and applying afterwards a linear transformation to the input HDR image.

Based on the works of Greg Ward, the creator of Radiance and pioneer of high dynamic range imaging, Viswanathan Kumaragurubaran in collaboration with Mehlika Inanici at the University of Washington, Seattle developed the hdrscope (Kumaragurubaran, 2013) software as a thesis project. Hdrscope is capable of performing HDR image processing and analysis for architectural lighting design and provides a Graphical User Interface for manipulating HDR Images including calibration methods. Hdrscope has been used ever since for teaching and various research projects (Inanici, 2017), and however good the documentation is, a lack of information is present about the difference in precision when doing luminance or illuminance calibration on HDR images.

Within hdrscope, luminance calibration can be performed by using a standard gray target in the scene, like the X-rite ColorChecker, and measure the luminance value with luminance meter. By capturing at the same time the HDR image set and measuring the luminance value on the reference card, it becomes possible to linearly calibrate the HDR image. Since it is not always possible to include a reference chart in the scene or have access to a luminance meter, the illuminance calibration can be used.

Illuminance calibration can be accomplished owing to the geometric relation between a lux meter and a fisheye lens. By placing the lux meter in front of the lens, with the sensor pointing towards the scene, lux measurements can be done before and after taking the photographs. Illuminance meters, sometimes called lux meters, employ cosine correction and are hemispherical measuring devices. Most commercially available fisheye lenses have either a circular or diagonal 180º field of view and perform a geometrical projection on the sensor plane. By transforming an equidistant fisheye projection from a circular fisheye lens to a hemispherical projection, the resulting image can be linearly calibrated with the mean lux value measured with the illuminance meter.

Mehlika Inanici(Inanici, 2017) and Alstan Jakubiec (Alstan Jakubiec, 2017) published several research papers about the use of HDR images for luminance mapping and capturing techniques, but in all of them it was not clear if images taken with a fisheye lens are needed for doing a luminance calibration, even though this is the only type of lens used for all the test cases. Generally, fisheye lenses are known for having several problems dealing with image quality like vignetting and chromatic aberrations, which can prove detrimental when combined with calibration factors, affecting the overall measurements.

Due to the lack of information on the matter, in the following sections I would like to raise several questions:

- are fisheye lenses necessary for creating a calibrated luminance map?
- Is it possible to use luminance or illuminance calibration with the same accuracy?

3.2 Capturing HDR images

The following list details the equipment used for all the tests in this chapter. The combination of camera and lenses has been selected specifically for these tests and the measuring equipment was borrowed from the Aalborg University Booking system.

Camera: Nikon D800 full frame DSLR
Lenses: Nikon 16mm f/2.8 AI-s, Sigma 8mm F3.5 Circular Fisheye, Sigma 50mm F1.4 Art
Luminance meter: Konica Minolta LS-150
Lux meter: Hagner Universal Photometer S2

As I mentioned before, a HDRI or a high dynamic range image is a combination of different photographs taken with different exposures, sometimes called brackets. The exposure is based on three parameters: shutter speed, aperture and ISO and any combination of the three parameters can provide a different exposure value, or EV. Bracketing the images based only on ISO increases the noise levels in the image and changing only the aperture creates a final image with varying degrees of depth of field and luminance. A better approach is to keep constant the ISO and aperture and vary only the shutter speed.

A fixed base, like a tripod, should be used at all times when taking HDR brackets. The longest and shortest shutter speed is dependent on the illuminance in the scene, and also the camera capabilities. Nikon D800 has an automatic bracketing system, but the maximum EV range is 9 bracketed shots with 1 EV step, which is insufficient for the test cases. Instead, the camera was tethered to a laptop and the SmartShooter software (Hart, 2017) was used to capture between 15 and 21 EV steps needed for all the tests. The same can be accomplished manually, but this would introduce motion in between different pictures and make the capturing process even longer.

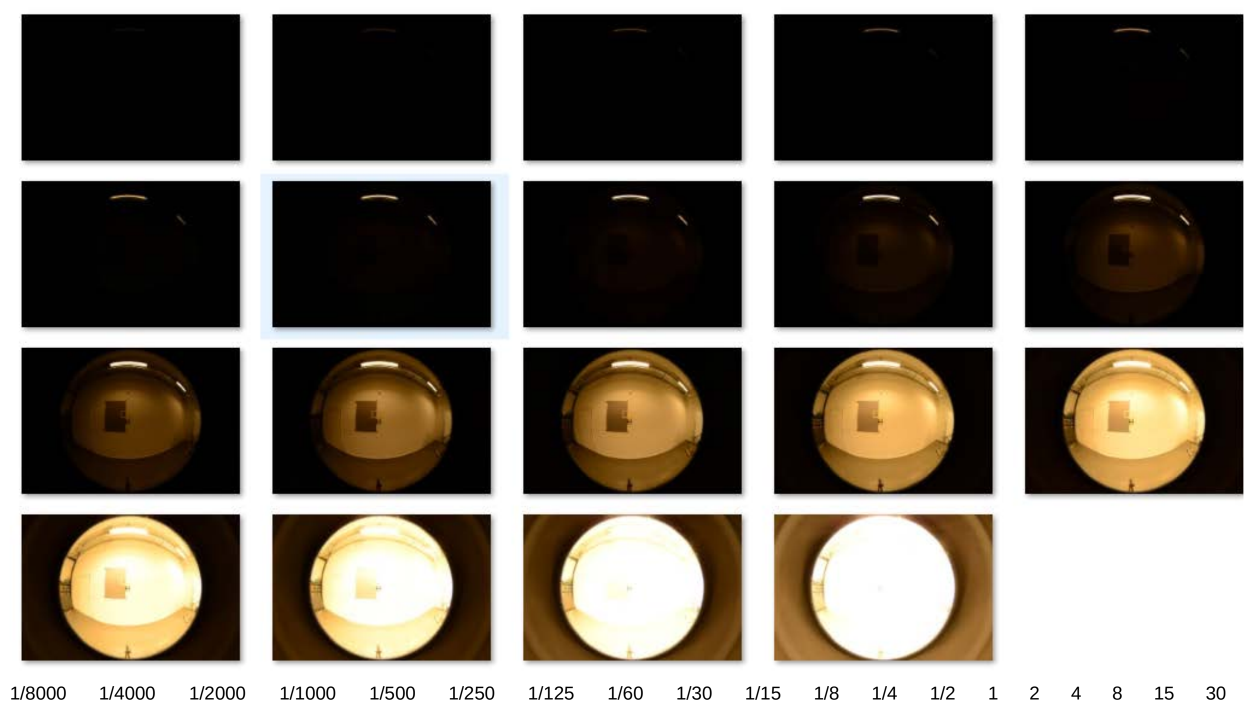

Figure 17 HDR set with 19 photographs 1 EV apart

On this example (Figure 17), the total time the camera was taking photographs was 65 seconds, which can be considered slow if the light in the environment is constantly changing. In this case any light measurements should be taken before and after the HDR set has been captured and the mean value should be used for calibration.

3.2.1 HDR Image Processing

The resulting images are post-processed in Photoshop to set the correct white balance and exported for Photosphere where the HDR final image is created. Based on a set of bracketed shots, Photosphere can infer the camera response curve and linearize the final HDR. This is a time saving feature, furthermore the information written in the file header is used later on by hdrscope to correctly apply transformations to the calibrated HDR image.

Depending on the type of calibration needed, two separate steps are taken in hdrscope. For luminance calibration, the area where the measurement was taken is selected and the measured value is specified for calibration. For illuminance calibration, the image needs to be cropped first to a perfect square around the image projected by the circular fisheye lens and then converted to a hemispherical projection. The last step is to specify the measured lux value and calibrate the converted HDR image.

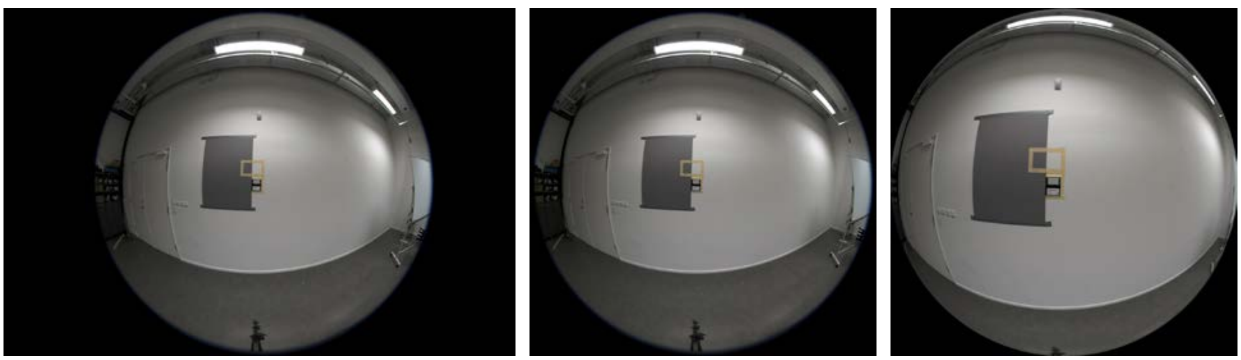

Original image Cropped Converted

Figure 18 Image stages for illuminance calibration

3.2.2 Calibration tests

Are fisheye lenses necessary for creating a calibrated luminance map?

The Light Lab room at the Aalborg University campus in Copenhagen was chosen for these tests as it contains a fully controllable environment for testing different lighting scenarios. Three different lighting schemes were created to do a comparison test.

For each test (Figure 19, 20, 21 and Table 5, 6, 7), the HDR image was calibrated using the luminance values measured on the reference chart and after calibration the pixels within the target areas were selected, analyzed and the mean luminance value was computed. All the measured surfaces are matte, lambertian materials that lack any glare that might interfere with the measurements. The maximum deviation encountered in all the tests was 11 cd/m2 for the spotlight test, where the light has a clear falloff within the target area.

Based on the documentation on how luminance calibration is being done and the results from the test, it is safe to assume that any type of lens can be used to capture an HDR image as long as the calibration is being done with values measured from a luminance meter.

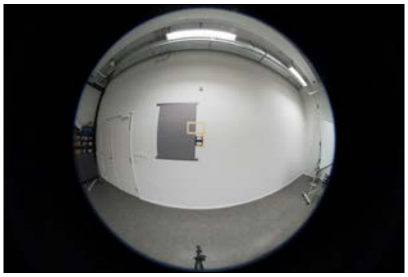
Sigma 8mm


Nikon 16mm

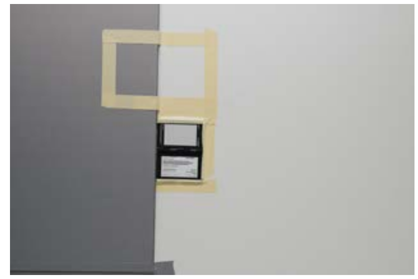
Sigma 50mm

cd/m2	LUMINANCE MEASUREMENT	AFTER LUMINANCE CALIBRATION		
LENS		Sigma 8mm	Nikon 16mm	Sigma 50mm
REFERENCE	20.8	20.8	20.77	20.71
WHITE WALL	33.86	32.8	33.24	33.15
GREY SURFACE	10.76	11.24	11.17	10.5

Figure 19 Ambient light from the ceiling; Table 5 Luminance calibration test


Sigma 8mm

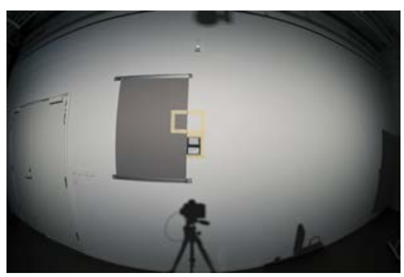
Nikon 16mm


Sigma 50mm

cd/m2	LUMINANCE MEASUREMENT	AFTER LUMINANCE CALIBRATION		
LENS		Sigma 8mm	Nikon 16mm	Sigma 50mm
REFERENCE	8.12	8.11	8.11	8.11
WHITE WALL	12.5	12.79	12.53	12.61
GREY SURFACE	4.1	4.09	4.21	3.91

Figure 20 Flood led lamp projecting from behind the camera; Table 6 Luminance calibration test

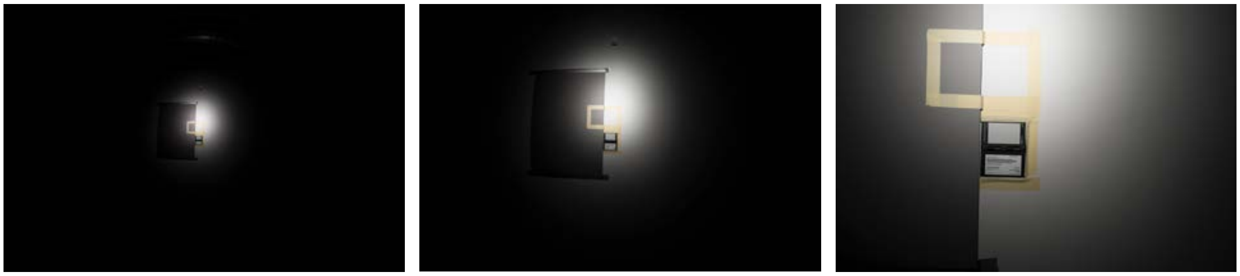

Sigma 8mm Nikon 16mm Sigma 50mm

cd/m2	LUMINANCE MEASUREMENT	AFTER LUMINANCE CALIBRATION		
LENS		Sigma 8mm	Nikon 16mm	Sigma 50mm
REFERENCE	204.5	203.88	202.23	201.21
WHITE WALL	749.4	743.49	754.54	745.03
GREY SURFACE	188.3	186.38	194.1	187.32

Figure 21 Spotlight aimed directly at the target area; Table 7 Luminance calibration test

Question 2 - Is it possible to use luminance or illuminance calibration with the same accuracy?

Illuminance calibration can only be used with a circular fisheye lens, and for the following tests all the photographs are taken using the Sigma 8mm lens. Different environments were chosen for the test, ranging from inside controlled lighting to outside sunny conditions. The reference card was placed in each scene and measurements were taken with both the luminance and the luminance meter.

In direct sunlight, the camera cannot capture the full range of illumination from the sun, since it would require an extremely fast shutter speed, however a ND or Neutral Density filter can be used to reduce the intensity of the light. Hoya is considered as one of leading manufacturers of optical equipment and their line of ND filters include a special type called Solas (Hoyafilterusa.com, 2016) that reduces the color casts occurring when photographing with increased exposure times. For these tests, the Hoya 77mm Solas IRND 3.0 Filter was used with a 10 Stop EV reduction factor, offering 0.098% light transmittance.

Changing lighting conditions affect the precision of illuminance calibration. In the first three examples with controlled environment for lighting (Figure 22, Table 8), the difference in between luminance and illuminance calibration are minimal, both scenarios producing comparable results. The following three examples with overcast sky present the same pattern except the example with the bridge. On this example, only one illuminance measurement was taken at the beginning of the image capturing process, and lighting conditions changed rapidly, due to rain. Failing to take a secondary illuminance measurement, or repeat the whole process at a different date yielded the difference in results.

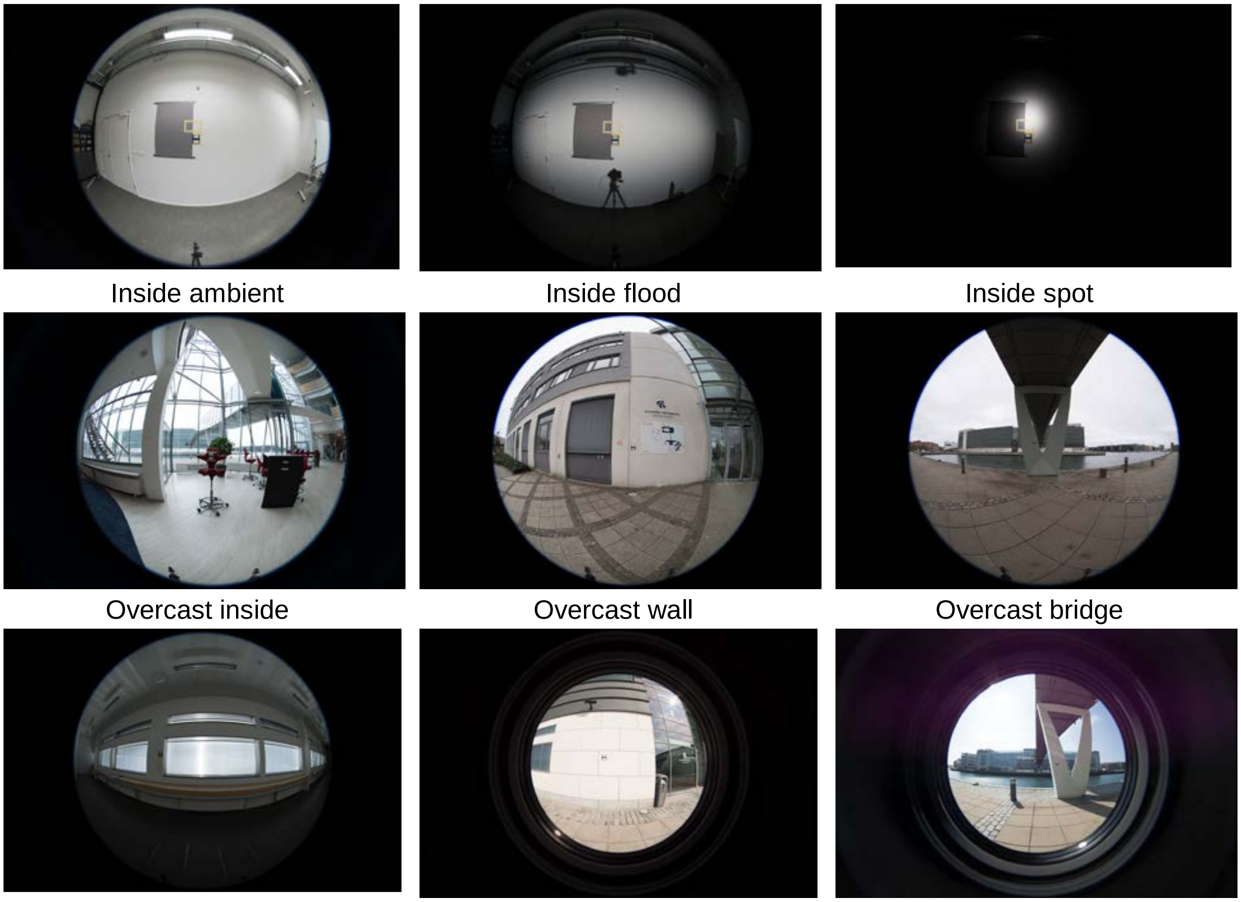

Inside ambient | Inside flood | Inside spot

Overcast inside | Overcast wall | Overcast bridge

Sunny inside | Sunny wall * | Sunny bridge *

Figure 22 HDR images for luminance – illuminance calibration tests

In the last example set, the test inside in front of the window with direct sunlight partially shielded by the blinds, the difference in between the two calibration methods is increasing slightly. The clear difference is provided in the cases where ND filter was used. The disadvantage in using a ND filter on the Sigma 8mm lens is the field of view reduction to 139º, hence the mathematical relation between the lens and the lux meter is being affected.

Analyzing the results (Figure 22, Table 8), I can conclude that in stable lighting conditions either calibration method can be used, with one warning that illuminance measurements should be taken before and after the capturing process.

Location	Illuminance measurement *Lm/m²*	Target	Luminance Measurement *Cd/m²*	After Luminance Calibration *Cd/m²*	After Illuminance Calibration *Cd/m²*
		Reference	20.8	20.8	20.81
Inside ambient	107	White wall	33.86	32.8	32.97
		Grey surface	10.76	11.24	11.29
		Reference	8.12	8.11	7.56
Inside flood	12	White wall	12.5	12.79	11.96
		Grey surface	4.1	4.09	3.82
		Reference	204.5	203.88	201.53
Inside spot	72	White wall	749.4	743.49	736.4
		Grey surface	188.3	186.38	184.87
		Reference	84	83.89	89.49
Overcast inside	1930	White wall surface	100	101.76	111.16
		Dark wood surface	7.5	8.2	8.72
		Reference	805	804.21	810.35
Overcast wall	1720	White wall surface	927	935.7	940.9
		Grey surface	146	154.7	156.81
Overcast bridge	4400	Reference	443	438.13	700.4
		Cement surface	418	393.5	645.92
Sunny inside	660	Reference	175.1	174.72	166.77
		Cement surface	239.5	237.2	223.89
Sunny wall *	28000	Reference	11960	11874	11587
		Cement surface	14180	13853	13497
Sunny bridge *	81100	Reference	2164	2161.6	2272.52
		Cement surface	1989	1981.35	2056.1

Table 8 Luminance – illuminance calibration tests

3.3 Capturing HDR images for IBL

Dedicated 360 cameras like Ricoh Theta or Samsung Gear 360 would seem like the ideal solution for this purpose, but these camera models have a low shutter speed and lack the option for using ND filters, making the process of capturing HDR images problematic in environments with high illuminance.

In order to create a spherical HDR image for use in Image-Based Lighting, a stitching software should be used to create the spherical panorama. There are several free and commercially available software options for creating panoramas, out of which PTGui is commonly considered one of the best (Ptgui.com, 2017). I will be using this for the next part of the process, but there are multiple options to choose a suitable software for different use cases (En.wikipedia.org, 2017).

Creating a 360-degree panorama involves taking multiple photographs from a fixed point while rotating the camera to capture images all around that point. The number of pictures needed is dependent on the focal length of the lens and the stitching software. For example, PTGui requires approximately 30% overlap in between the images for a successful panorama. Secondly, to avoid parallax errors it is recommended to use an adapter called a panorama head that maintains the lens no-parallax point at the center of the tripod. This point is based at the center of the lens's entrance pupil and in the photography community, this special point is often called the nodal point.

To exemplify the process (Figure 23), I took 8 photographs with a Nikon 16mm lens covering both the zenith and the nadir, or the top and bottom of the scene. The RAW images were imported to PTGui for stitching and even if the images had a good coverage, it still required manual adjustments through control points.

Figure 23 Nikon D800 with Nikon 16mm lens panorama test with 8 images

The total capture time, without bracketing, was approximately three minutes, which is too long for capturing accurate HDR images in environments with changing light. Furthermore, the use of the panorama head is not optimal since the resulting panorama had anyway a missing area at the bottom that had to be cleaned in Photoshop (Adobe.com, 2017).

Following this initial assessment, several different options were tested (Figure 24) to reduce the capture time, but also provide a usable HDR panorama for De-lighting purposes. Firstly, Sigma 8mm lens was used since it has the largest field of view and can also be useful for situations where illuminance calibration is needed. Secondly, the number of individual shots was reduced to three keeping with the recommendations for image overlap and then the pictures were taken with a panorama head and a normal ball head to asses the difference.

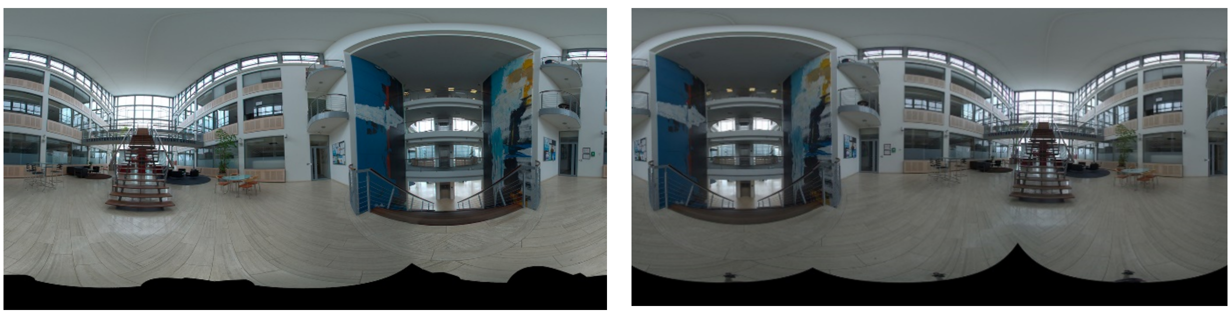

Panorama head plus nadir shot | Ballhead only three shots

Figure 24 Nikon D800 with Sigma 8mm panorama test with two different tripod heads

Lowering the number of individual orientations increases the capturing speed, without any detriment in the quality of the final panorama. In the first image a panorama head was used to capture three different orientations and also the nadir, but the resulted panorama still had the bottom 10% missing. Furthermore, comparing to the second option where only three images were taken using a normal ball head, the absent area is similar and still requires further processing in a digital painting application.

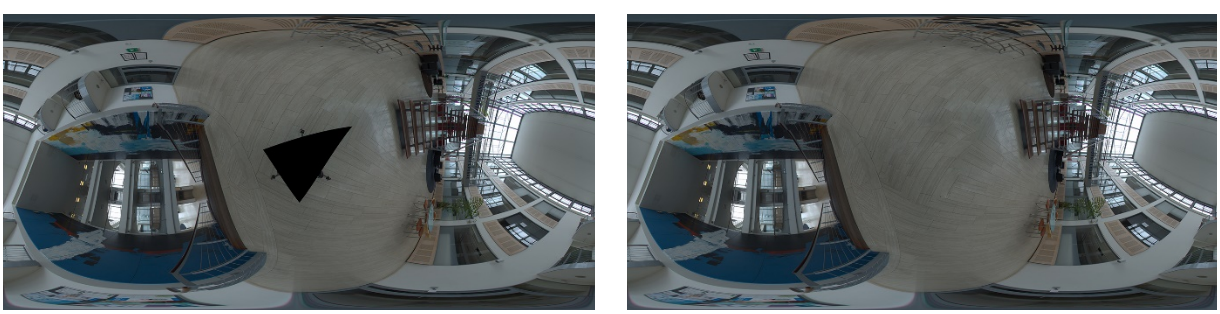

Figure 25 Nikon D800 with Sigma 8 mm panorama, retouching the nadir in Photoshop

The final workflow includes taking three photographs with a Nikon D800 and the Sigma 8mm lens. It is important to capture the zenith, so the camera and lens can be place with a small 5-10º angle orienting upwards on a normal ball head. To cleanup the resulted spherical panorama image, the Photoshop plugin Flexify (Flamingpear.com, 2017) can be used to change the initial orientation of the panorama image and retouch the nadir area (Figure 25).

Figure 26 Final panorama

The final 360 panorama image (Figure 26) required less than one minute to capture, involved using a camera and lens that can take circular fisheye images suitable for calibration and with a total resolution of 9000/5000 pixels is suitable for Image-Based Lighting.

3.4 Use case – the Thinker statue

During the general capturing process for the Thinker statue I also took the photographs needed for the HDR de-lighting procedure. Based on the same workflow discussed previously, three bracketed sets of 15 images each were taken in front of the statue, and illuminance measurements were taken twice for each set. I also measured the luminance value on the X-rite ColorChecker card for data safety reasons and to compare the results at the end.

Each set of photographs was transformed into an HDR image using Photosphere and then imported to hdrscope for calibration. Running the initial illuminance calibration on each set, the images were modified by changing the projection type and were unusable for stitching. To fix this, small areas that contained a matte material were selected and the luminance value noted. A second set of calibration was performed on the initial HRD image using the luminance calibration method (Figure 27).

Figure 27 luminance calibrated HDR images

These pictures were than imported to PTGui for stitching and later retouched the nadir in Adobe Photoshop. The resulted HDR panorama was imported back to hdrscope for verification (Figure 28). The average deviation of luminance values in between the initial HDR image and the final 360 HDR panorama was less than 1%, which in my opinion establishes that the workflow is nondestructive and suited for calibrating images for IBL.

Figure 28 Calibrated HDR panorama

In order to bake the lighting information to a texture, Autodesk Maya with the Arnold renderer (Autodesk.com, 2017) was selected. There are many 3D software applications that can be used for this process, but the important element when choosing one or the other is the bit depth of the resulted texture. HDR images are 32-bit image files and Arnold renderer can bake 32-bit EXR file, thus maintaining a nondestructive workflow.

The 3D model of the statue with the UV map was imported to Maya together with the HDR panorama and aligned so that the 3D model of the statue is facing the same orientation as the real-world object (Figure 29). The Arnold Render Selection to Texture utility was used to convert the illumination to single texture, portraying how the material of the object was affected by the available light in the environment at the time of capture (Figure 29).

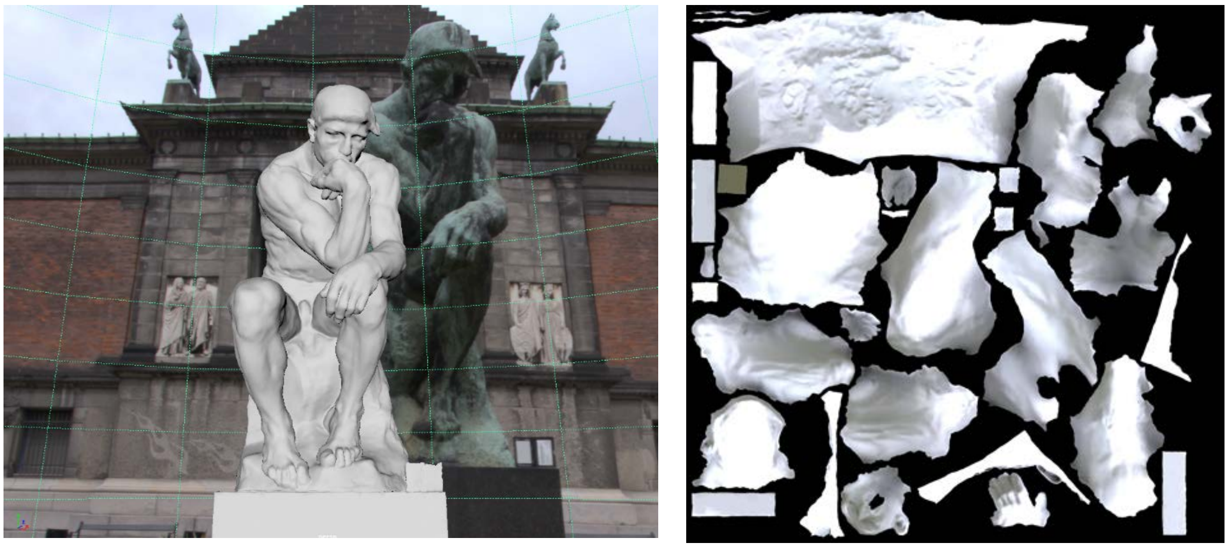

Figure 29 Object position surrounded by the HDR panorama and baked luminance

The final step in the de-lighting process is done in Photoshop. The initial texture exported from Reality Capture is imported to Photoshop and converted to 32-bit. Also, the illumination texture is imported as two exposure masks in Photoshop, one for the shadows and one for the highlights. Afterwards, the exposure sliders are used opposite of each other to balance out the final texture (Figure 30).

Currently, the results are useful if there is less than 4 stops difference in between the exposure levels for shadows and highlights. Using an 8-bit image file for the initial diffuse texture is insufficient to accurately sample a continuous progression of color tone until artifacts like posterization start appearing. The current limitation is due to Reality Capture's inability to export 16 or 32-bit textures at the time of writing this, though taking into account their current plans to support this feature by the end of September 2017, this method should prove even more useful for scenes with higher dynamic range.

To assess the difference in between the two textures, the 3D model of the statue was placed in a scene with one directional light pointing towards the front of the statue (Figure 31). The model on the left has no textures, just the base geometry, the model in the middle used the original diffuse texture and the model on the right has the de-lit version of the diffuse albedo. For this particular test, I consider the procedure to be successful, as it portrays closely how the original material of the statue is behaving in accordance to changing the light conditions and it is ready to be placed in a scene with focus on dynamic light and possible interactivity.

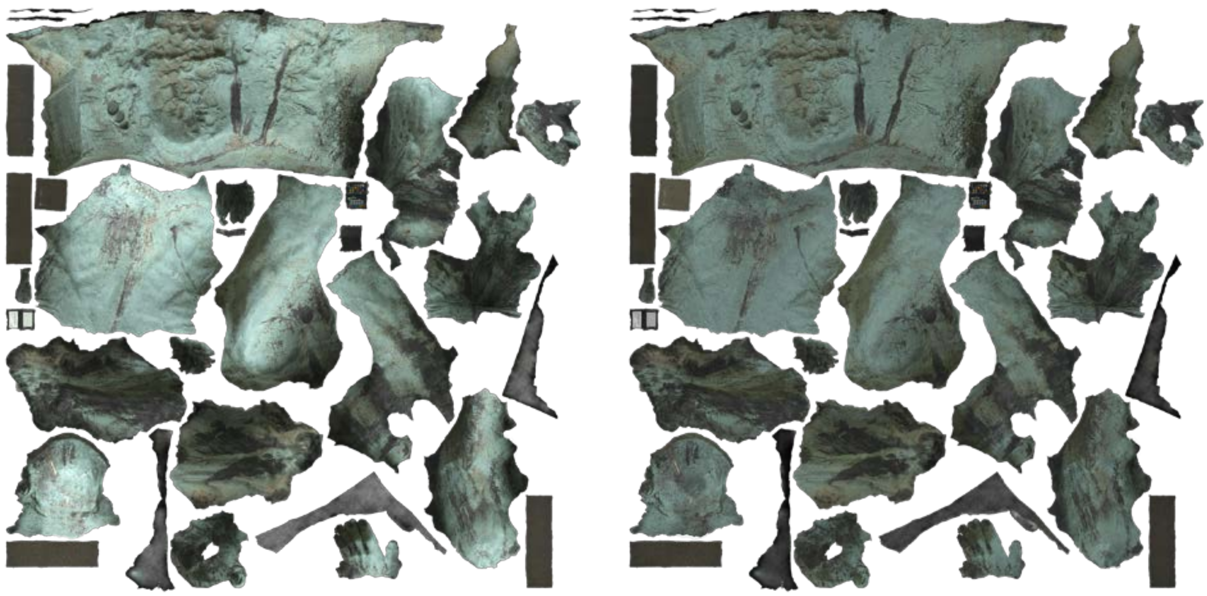

Figure 30 Texture with and without baked light information

Figure 31 Render of the object without texture, with original texture and with the de-lit texture

Though the HDR de-lighting procedure can yield good results by removing the light information baked in the diffuse texture, it does not always work properly. If any of the steps in the workflow are done incorrect, then the resulted assets will behave erroneous in environments with dynamic light.

3.5 Capturing only the albedo

As I mentioned before in the PBR workflow introduction, the diffuse albedo represents the base color of the object independent of viewpoint. Also, this map should be void of any lighting information, be it specular highlights or shadows.

In photography, polarizing filters can decrease reflections or specular highlights and are one of the few filters that cannot be replicated by any photo editing software. They are placed in front of the camera lens and work by filtering out the sunlight that is directly reflected from plane surfaces like water or windows. The angle by which the light is filtered out is controlled by physically rotating the filter in front of the lens (Figure 32).

CPL Filter 0º CPL Filter 45º CPL Filter 90º

Figure 32 Circular polarizing filter test

In general, objects in direct light appear more desaturated because of the specular light diffusion encountered on the surface of the material. By using a polarizing filter, the specular diffused light is filtered out and color saturation is increased. This can be considered the base for the disuse albedo map, because it preserves the color information and the specular component is removed. Furthering the process, by applying a subtraction filter between the polarized and the unpolarized images we can separate the specular component and thus create the base for the specular map (Figure 33).

However, light coming from the sun or most electric light sources is not polarized, and in the example with the car, polarization is achieved because of Brewster’s angle. When light strikes a reflective surface at a certain angle called the polarizing angle or Brewster’s angle, the reflected light becomes completely polarized (Encyclopedia Britannica, 2017).

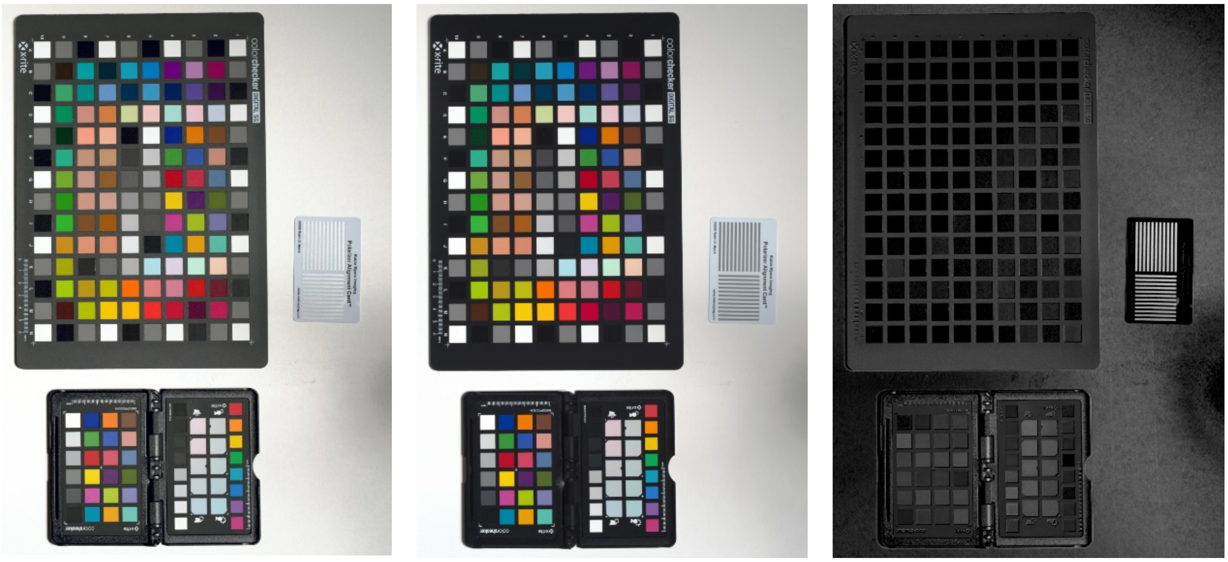

Figure 33 Unpolarized, polarized and generated specular difference

Since this light effect is only valid for a few instances, a different approach needs to be taken for consistent results. To ensure that only polarized light is present in the scene, a polarizing filter needs to mounted on an electric light source as well. Polarizing the light has a detrimental effect that the available amount of light is reduced to generally 40% depending on the transmittance of the filter. Flashes are the most viable option because they can emit light with high intensity that can overpower the ambient light even when polarizing filters are mounted on the flash.

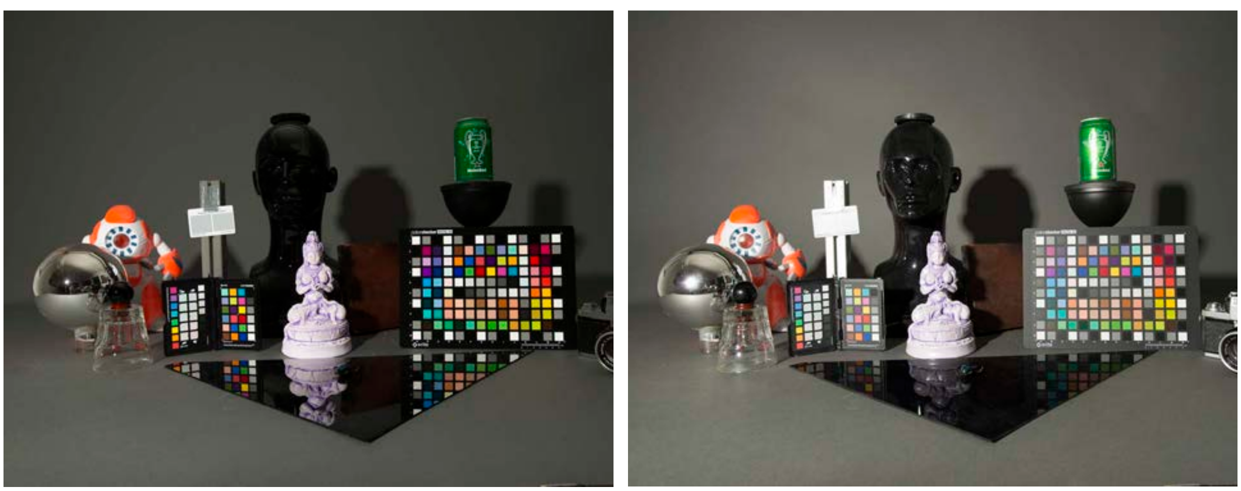

Figure 34 Cross polarization setup is polarized and unpolarized light

The technique of combining two polarizing filters, one on the light source and one on the lens positioned perpendicular to each other is called cross-polarization (Figure 34). One important limitation to using cross polarization is the fact that the light source and the camera must be on the same plane, otherwise the polarization angle switches and the object becomes unpolarized.

To test this approach in an uncontrolled environment, a Nikon D800 camera with a Sigma 50mm lens was used with a circular polarizing filter mounted on the lens. Two battery powered flashes with polarization filters were positioned side-by-side to the camera with their orientation angle set to 45^{o} toward the object and polarizing filters rotated for cross-polarization (Figure 35).

Figure 35 Portable cross-polarization setup with 1 camera and 2 flash light sources

In the field, it became obvious that this setup is cumbersome, weighing without a tripod approximately 8 kg (Figure 35). Nevertheless, the object was captured using only the images with polarized light coming from the flashes, and later reconstructed in Reality Capture. In total, the capturing process took approximately 1.5 hours for 560 images.

Taking photographs with a cross-polarized system has clear advantages in contexts where reducing the specularity of an object is important, but the setup used for the test can be considered unwieldy for long term use. The flashes need to be powerful enough to overpower the light in the scene, but also as lightweight as possible to permit capturing from different angles, especially the top parts. Secondly, as the batteries for the flashes wear out and the recycle time increases, the units are prone to misfires, making the whole process impractical.

Using the cross-polarization technique should not be viewed as a replacement for the de-lighting technique even though both offer the same result at the end, a diffuse texture without any baked lighting information. The context in which either technique can be used is important. For example, the de-lighting technique is best used for large objects or even buildings, where it is easier to capture an HDR panorama than move around with camera and light equipment.

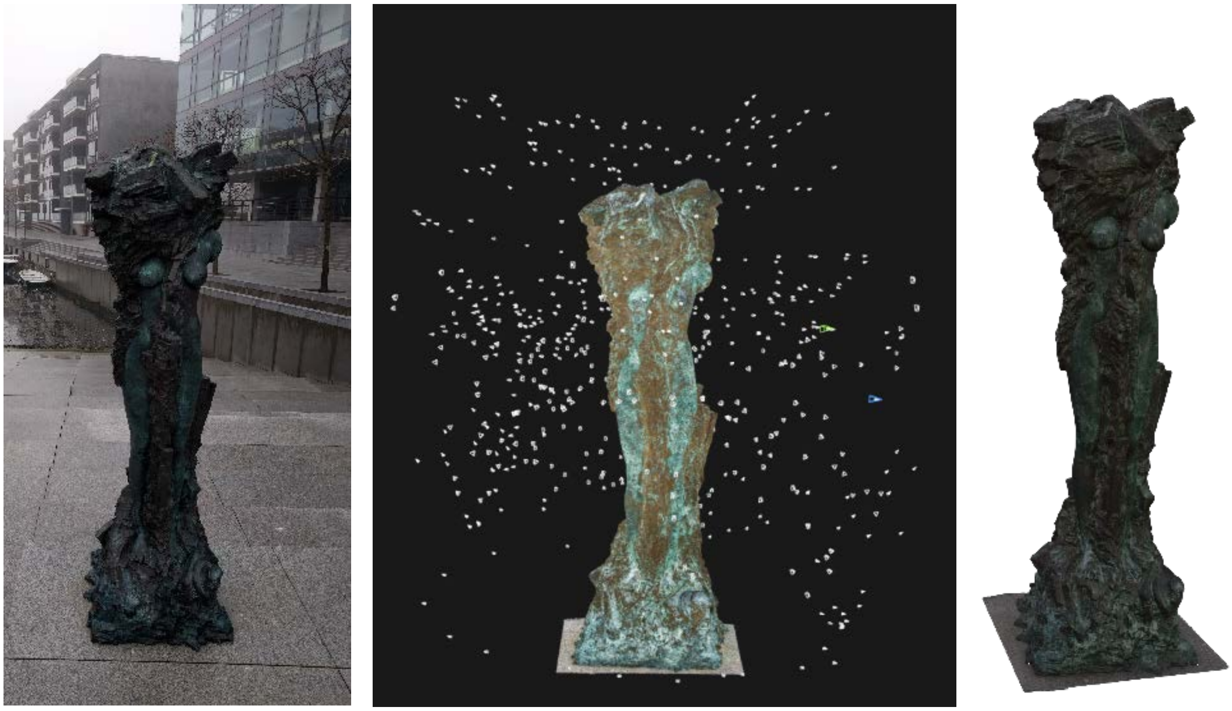

Photograph 3D reconstruction with camera locations Render

Figure 36 Statue reconstructed using only polarized light images

Furthermore, there are situations where the object dramatically changes the way it looks under cross-polarized light, like gold which becomes a dull brown color (Figure 36). However, there are cases where the geometry of the object cannot be captured unless the specularity is reduced, and this is where cross-polarization technique is better used, for capturing objects with high reflectance.

To sum up, both techniques have advantages and disadvantages, and even though each negate the use of the other, the choice depends on the context and the limitations of the environment.

Chapter 4. The road to automatic capture

Cross-polarization has proved to be a good aid tool for photographing dielectrics that are glossy or present high specularity. An important limitation is that the light source and the camera must be on the same plane. Since photogrammetry reconstruction works better when there is a clear movement from picture to picture, a setup that allows to move both the camera and lights together is needed.

A good way to insure optimum image capture is to employ a robotic system that can take all the photos needed in a programmable and precise fashion. A brief analysis into the robotics market (Kuka.com, 2017) brought forth some ideas, but also a clear statement: the entry prices were tenfold the budget, and systems are proprietary, without the ability to modify them for different tasks without service from the manufacturing company.

To help advance the understanding of digital reconstruction and act as a hardware platform for testing varying workflows and technologies a robotic system was designed, built and continually improved and tested.

The following criteria were set in order to help as a guideline during the design process:

- Modular, easy to modify or improve
- Precision and repeatability
- Off the shelf electronics and actuators
- Optimum coverage with fast capture
- Enable cross-polarization capture
- Extensible for mixed technologies

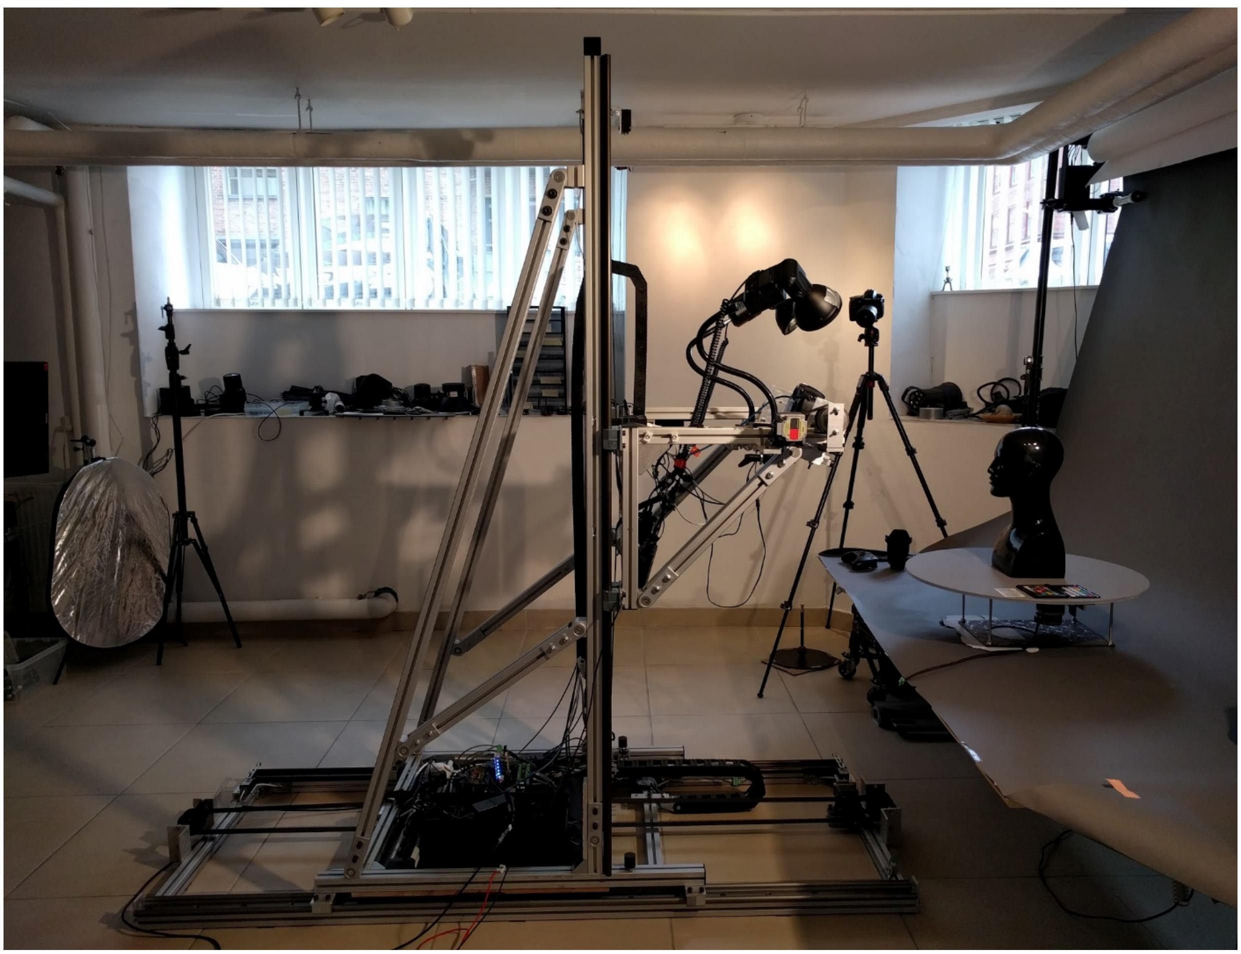

Figure 37 Side view of the camera rig

4.1 Mechanical and electronic design

Since a robotic arm with a high payload was out of the question, different industrial approaches for linear movement were analyzed and the gantry was chosen as the base system. A classic 3D printer, CNC machine or laser cutter is a type of gantry – it allows for precise positioning on one or multiple axis by converting rotational local movement from an actuator to linear movement on a specific axis.

The robotic system or camera rig (Figure 37) as it called now is a 2-piece system having a turntable and tiltable arm that could drive to a specified location and position the camera on any angle relative to the object. Based on simple calculation of weight for various camera, optics and light sources, the total weight of the equipment could reach 20 kg. Custom 3D printed parts were not an option, but a more industrial approach was needed. It had to be durable enough to hold heavy camera equipment, but also light enough to be moved with ease by small electrical motors. Aluminium, in general, is a good starting place, since it is both light and durable and aluminium profiles offer a good balance between weight and structural strength. (Stockist, 2017)

The second step is figuring out the actuating mechanism and the linear rail system (Figure 38). Several options were analyzed for usability, but after several meetings with a representative from the company, Igus products were chosen. The main advantage over other techniques, like recirculating ball bearings, is that their linear systems do not require lubrication, are durable enough for the task at hand and also stock everything needed to start building the prototype. (Igus.eu, 2017)

The meetings also helped with the decision between belt and pulley pairs or leadscrew modules. Belt and pulley systems are less prone to stalling if misalignments are present in the overall implementation. Considering the fact that this system is built in the premises of the AAU University in Copenhagen without the use of advanced CNC machined parts, it is inevitable for minor errors to occur. What was more important is that the software control system could account and adjust for this.

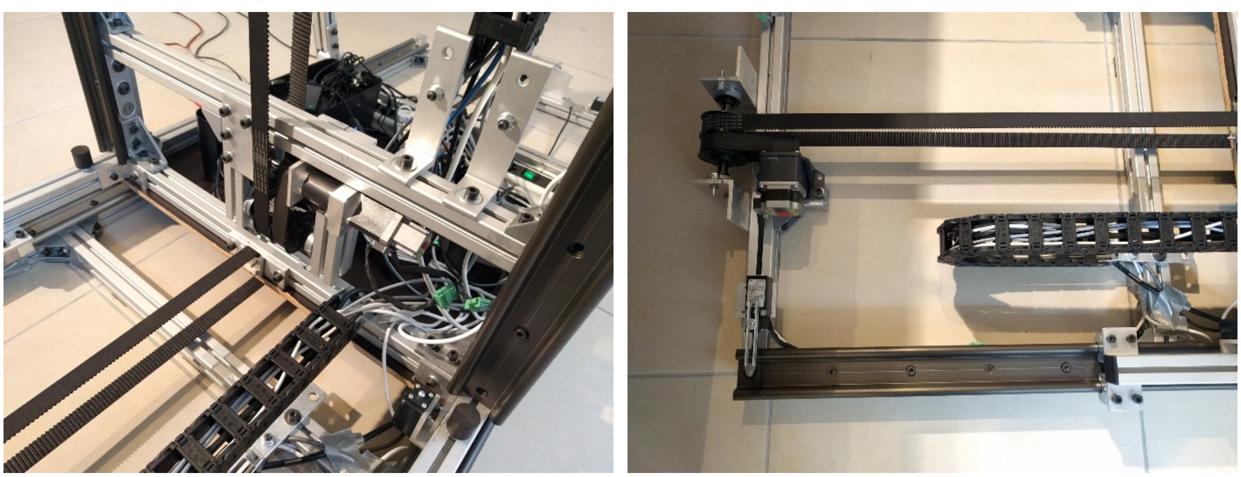

Figure 38 Close-up showcasing the stepper motors, linear rail and belt driving system

With the mechanical part done, the attention switched to the motors, electronics and software control system. One of the main design criteria being precision, a closed loop servo motor system would be ideal(Orientalmotor.com, 2017), however the cost was too high for the budget. The decision was made to choose an open loop system and find another way to insure precision, which came in the way of stepper motors (Figure 38). Stepper motors feature high torque and low vibration at low-speeds, ideal for applications requiring quick positioning for short distances (Orientalmotor.com, 2017). Also, they offer the possibility to request constant updated on position, making them suitable for this application.

During the studies, we were presented with a suite of various approaches and platforms suitable for prototyping. Arduino would be the classical choice for fast prototyping projects, but it lacks the rugged approach needed for this prototype, since the motors are quite powerful and I could not find an Arduino product that matched the requirements. Instead, TinkerForge products were used to build the complete set of hardware controllers needed to successfully operate the prototype. (Tinkerforge.com, 2017)

TinkerForge products, much like Arduino, are based on a flexible design that involves the use of modules with different tasks that can be put together in stacks to form complex systems. Also, they offer the possibility for wireless communication between different modules, task suited perfectly for this prototype.

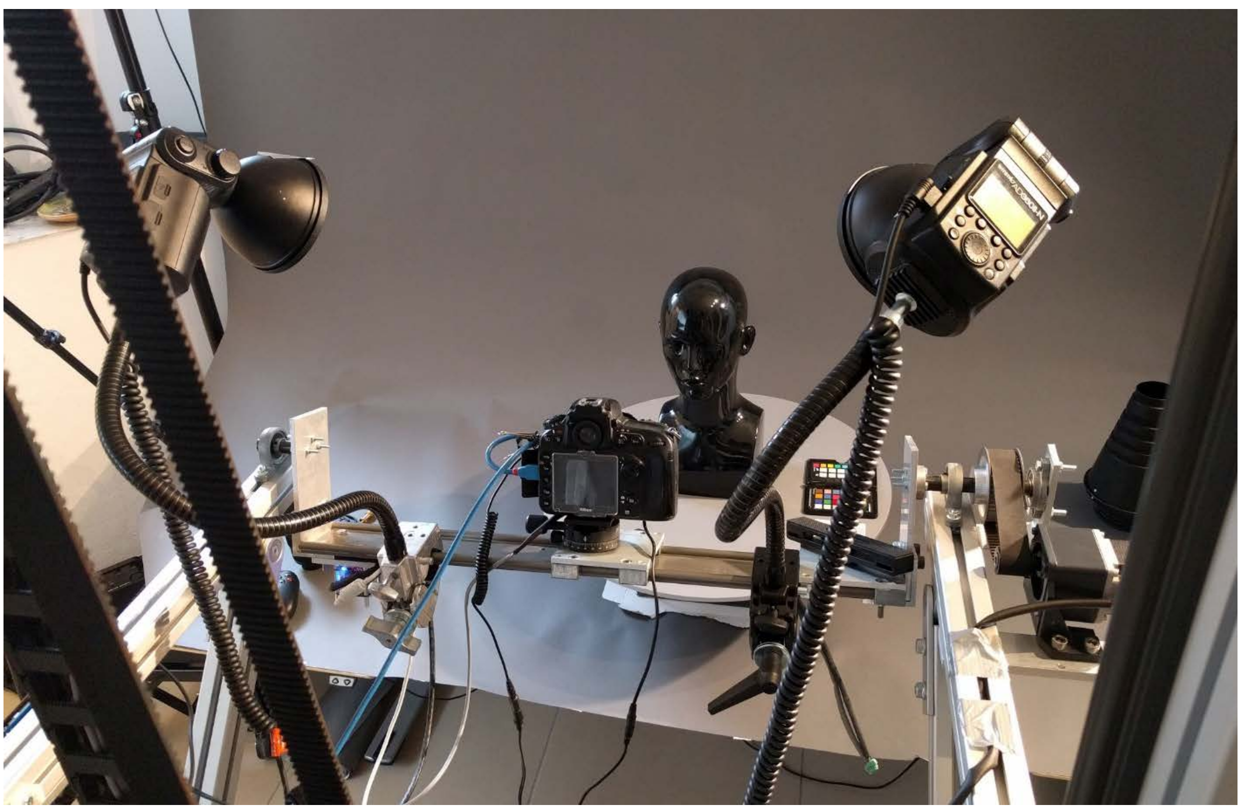

Figure 39 Inside close-up view towards the camera position

The total height of the camera rig prototype is 200 cm and the total capture arc is 170 cm with an overhang distance of 50 cm. Depending on the size of the object and required precision, the maximum object bounding box is 170 cm height / 150 cm width. For such a large object, being able to control both the turntable and the camera rig offers the needed coverage in a precise manner.

Speed of capture is important, especially if the rig is used in an environment that requires a large amount of individual captures, like a university or a museum. Since the system is capable to operate with 6 DSLR cameras simultaneously, considering also the flash recycle time or the shutter speed for longer exposures, it is safe to assume 1 image capture per second or 1fps. During the tests, it was observed that most objects require between 300 and 900 photos for an accurate reconstruction.

The rig is also placed on wheels to increase usability and also to account for objects that cannot be placed on a turntable. This way, positions can be recorded and capturing can be done on much larger objects, though this workflow would probably be used in special cases.

4.2 Control software

A custom software was developed for controlling the hardware components of the system in a flexible and safe manner. The main functionality of the software includes the communication with various Tinkerforge building blocks by using their Python API and controlling the behaviour of the individual hardware elements like stepper motors, safety switched, camera triggers, etc.

In total, the rig utilizes four stepper motors that correspond to four motion directions/axis:
X - linear horizontal movement of the system
Y - linear vertical movement
R - rotation of the tilting element on which the camera equipment is placed
T - rotation of the turntable on which the object is placed

In order to translate the camera position in space to a metric value, the software is able to run the rig through a calibration stage. By knowing the exact physical distance the system can travel on each axis, we can translate the amount of steps it took for each motor to move between start and end motion points. When knowing the exact distance to the object and the camera equipment at hand, the capturing process becomes much more precise and versatile.

The interface of the software (Figure X) provides multiple functions for setting up desired camera positions and angles in space and the amount of needed coverage. Currently, initial setup of the system consists of manually driving each axis on the desired position (by using corresponding buttons or the programmed wireless controller.) Each key position can then be recorded by registering four corresponding values of each motor. Moreover, several other options for capture can be defined:

- Amount of turntable stops per position defines how many images will be captured per key position. In case we define 36 stops and 360 degree turntable limit per camera position, the system would move the object taking a photo every 10 degrees. When the sequence is finished, the rig would drive to the next camera position.

- Turning on the polarization option activates the double-capture per position - once with the camera polarization filter active and once disabled. Capturing the object with this option enabled would result in twice the amount of images captured. The difference in between two datasets would be the amount of reflections and overall specularity of the image. As mentioned earlier, the use of cross polarization can reduce the amount of light reaching the camera sensor by up to 40% depending on the filter in use. In order to avoid the difference in exposure between the image sets, strobe flash lights are switched to a TTL* mode instead of using the manual power output.

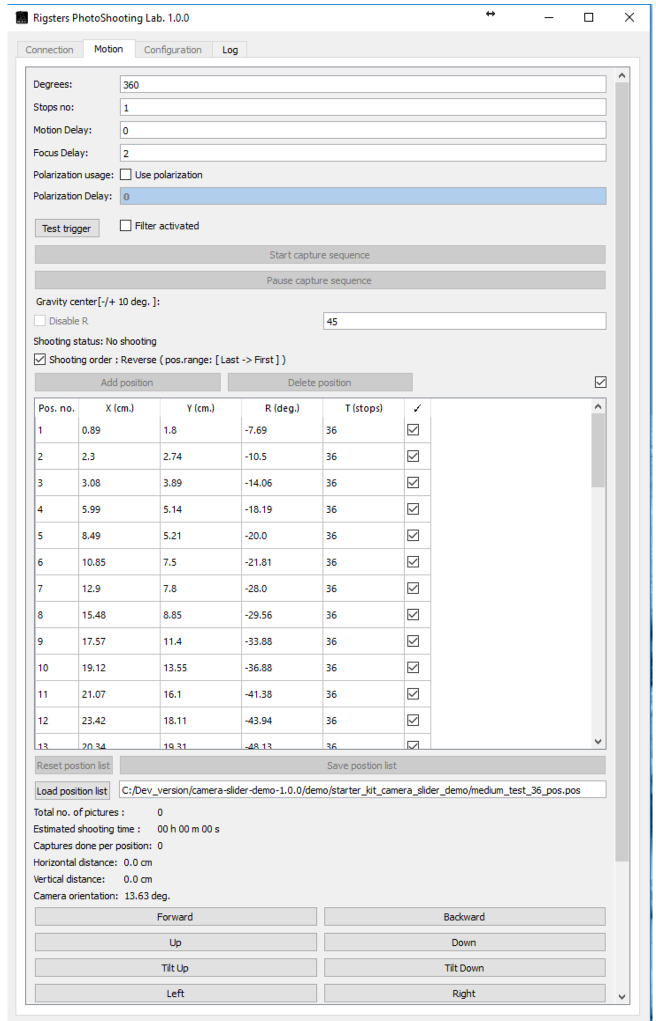

Fig 40 Camera rig custom control software GUI

- Various delay options are used in certain scenarios depending on the camera setup and focusing modes. Generally, by giving the camera rig a small delay before the capture would eliminate possible focusing and motion blur issues.

Though the initial camera setup can be still a troublesome, time-consuming process for an operator, the possibility of recording camera position lists and creating a library of relevant presets is the current remedy.

4.3 Camera rig use cases

The main purpose to build the capture rig is to act as a hardware platform for testing varying workflows and techniques. In the following examples, I would like to showcase several capturing sequences that would either be very time consuming or impossible to do just by using the handheld approach or even a tripod.

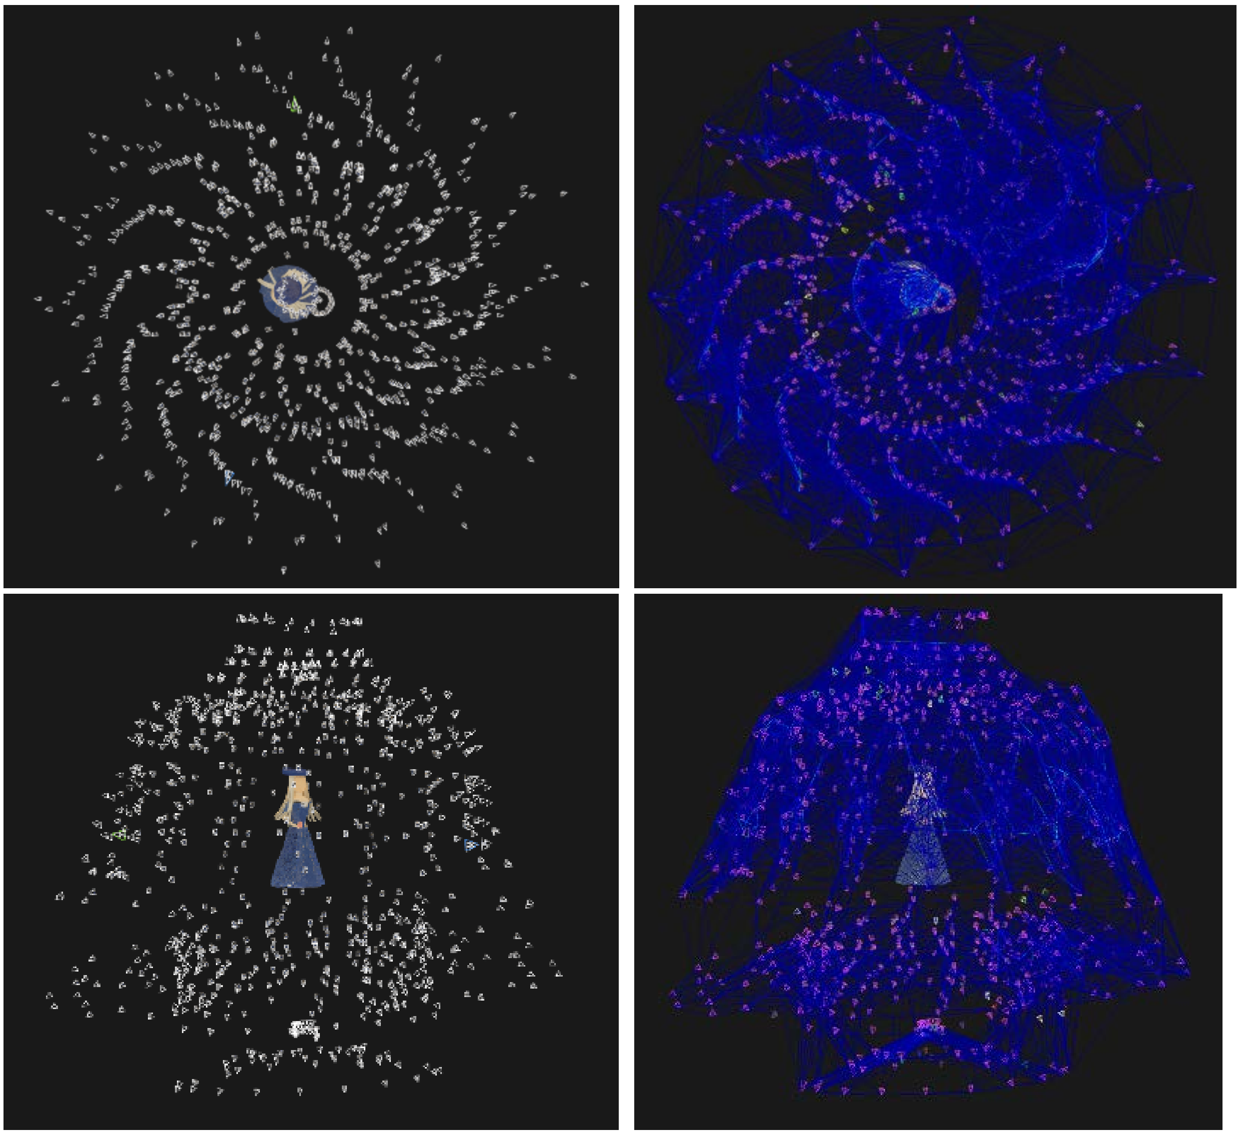

Figure 41 Top and side view from Reality Capture showcasing intense camera coverage

This example represents the possibilities for intense coverage in order to capture a challenging object. In total 960 images were taken using one camera, in approximately 35 minutes. The object was first positioned on the turntable, and 20 different elevation angles were set with various stops per turn, between 18 and 54 depending on the complexity of that specific area. After this first capture set was done, the object was flipped and another similar set was captured for the underside.

However, what is most important about this camera rig is the ability to capture images with polarized and unpolarized light. As discussed in the previous chapters, there are several ways to create or derive a diffuse albedo map, but the specular map obtained by the difference in between polarized and unpolarized images can only be done using a tripod or a fixed base (Figure 42). The camera rig allows the capture of these images by using an active polarizing filter connected to a custom relay that changes the polarization state in under 1/10th of a second.

Photograph with unpolarized light

Photograph with polarized light

Diffuse Albedo map

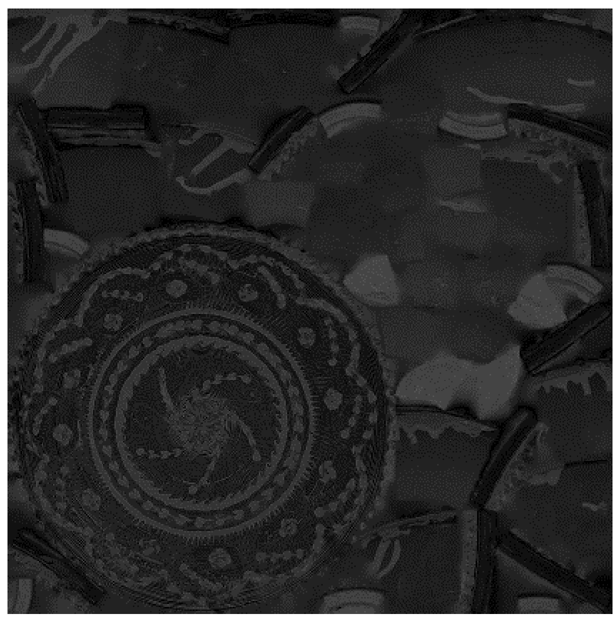

Specular map

Figure 42 Showcasing the Diffuse Albedo and Specular map created using the camera rig

A lighting designers work does not always center around lighting schemes for office environments or interactive installations, but it can take various artistic and technical purposes. For example, the same workflow was utilized for scanning a dozen of leaves that are currently used within a research project in advanced visual recognition systems at Aalborg University in

Copenhagen. My work for this project involved capturing different parts of the plants, like leaves and stems and create accurate materials and textures used for simulating how the plant appears under varying degrees of illumination.

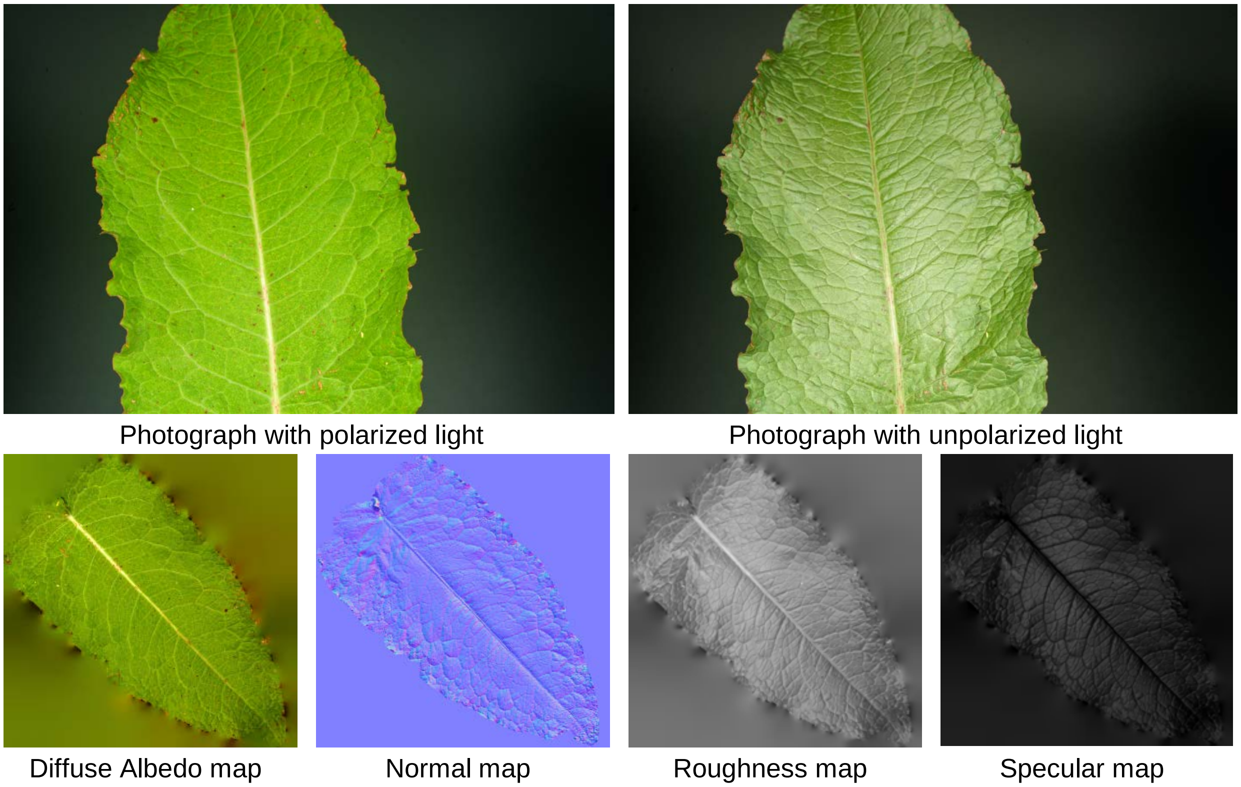

Figure 43 Showcasing the texture maps created using the camera rig

4.4 Further work

There are several further development concepts in the works for the current prototype, that could improve the overall workflow. Firstly, a ZED stereo camera (Stereolabs.com, 2017) will be placed on the rig to act as a real-time depth sensor and constantly scan the environment for changes, increasing the security of the rig, but also allowing the system to pre-scan an object and make recommendations for capturing strategies.

Secondly, a new lighting system with 64 leds will be mounted on the rig to permit the capture of roughness maps. This setup is based on the principles of RTI or reflectance transformation imaging where the roughness of a material is deduced by analyzing several images of the object with various lighting gradients.

Thirdly, the control software will be ported to C# in Unity to increase the interoperability with the ZED stereo camera and the future Reality Capture SDK (YouTube, 2017). This move is important, as it will permit the use of one umbrella software to both capture and process the images.

Chapter 5. Conclusion

The aim of this paper was to explore different methods for capturing 3D models using photogrammetry and the steps required to use these models for lighting visualization purposes in real-time time engines like Unity. One of the reasons photogrammetry should be considered by lighting designers is the ability to capture both the geometry and the texture of the object.

Having the object correctly react to dynamic light in a physically-based rendering engine is still the most problematic area as currently there are no simple methods for creating all the necessary maps. De-lighting techniques using HDR images and cross-polarization advance the process, however there are still parts in the procedure that require artistic approach.

An automated system can help mitigate some of the issues by offering the possibility for capturing images used to generate specular maps. A secondary advantage is the increase speed of capture which can prove very helpful if the lighting designer is working with a large number of objects or materials.

Further work on the automated system can help reduce the number of individual processing steps and possibly offer a one-click solution for the overall process.

References

Adobe.com. (2017). *Adobe Photoshop CC*. [online] Available at: http://www.adobe.com/dk/products/photoshop.html [Accessed 7 Apr. 2017].

Agisoft.com. (2017). *Agisoft.com - User Manuals*. [online] Available at: http://www.agisoft.com/downloads/user-manuals/ [Accessed 1 Sep. 2016].

allegorithmic. (2016). *PBR Guide*. [online] Available at: https://www.allegorithmic.com/pbr-guide [Accessed 23 May 2016].

Alstan Jakubiec. (2017). *J. Alstan Jakubiec - Architecture and Sustainable Design (ASD)*. [online] Available at: https://asd.sutd.edu.sg/people/faculty/j-alstan-jakubiec [Accessed 8 Jun. 2017].

Autodesk Remake - Getting Started Guide. (2017). [ebook] Autodesk, pp.48-53. Available at: https://remake.autodesk.com/assets/Autodesk_ReMake_Guide_01.pdf [Accessed 4 Sep. 2016].

Autodesk.com. (2017). *Maya Autodesk*. [online] Available at: https://www.autodesk.com/products/maya/overview [Accessed 12 Jun. 2017].

Bentley Communities. (2017). *ContextCapture Guide for Photo Acquisition*. [online] Available at: https://communities.bentley.com/products/3d_imaging_and_point_cloud_software/m/mediagallery/271352 [Accessed 1 Sep. 2016].

Cambridgeincolour.com. (2017). *Photography Tools*. [online] Available at: http://www.cambridgeincolour.com/photography-tools.htm [Accessed 1 Aug. 2015].

Capturingreality.com. (2017). *CapturingReality.com - Tutorials*. [online] Available at: https://www.capturingreality.com/forum/viewforum.php?f=6&sid=7a5db3e1888cef3e1aae1a0f33e7642d [Accessed 1 Sep. 2016].

D&M Imaging. (2017). *Bit Depth*. [online] Available at: http://dmimaging.net/8-bit-vs-16-bit-images/ [Accessed 13 Jun. 2017].

Debevec, P. (2004). The Parthenon. *ACM SIGGRAPH 2004 Computer animation festival on - SIGGRAPH '04*.

Docs.unity3d.com. (2017). *Unity Documentation*. [online] Available at: https://docs.unity3d.com/Manual/StandardShaderMaterialParameterNormalMap.html [Accessed 12 Mar. 2017].

Dxomark.com. (2017). *DxOMark*. [online] Available at: https://www.dxomark.com/ [Accessed 5 Apr. 2017].

En.wikipedia.org. (2017). *Brewster's angle*. [online] Available at: https://en.wikipedia.org/wiki/Brewster%27s_angle [Accessed 26 Feb. 2017].

En.wikipedia.org. (2017). *Comparison of photogrammetry software*. [online] Available at: https://en.wikipedia.org/wiki/Comparison_of_photogrammetry_software [Accessed 12 Jan. 2016].

En.wikipedia.org. (2017). *Hugin (software)*. [online] Available at: https://en.wikipedia.org/wiki/Hugin_(software) [Accessed 24 Jul. 2016].

Encyclopedia Britannica. (2017). *Sir David Brewster | Scottish physicist*. [online] Available at: https://www.britannica.com/biography/David-Brewster#ref261121 [Accessed 26 Feb. 2017].

Flamingpear.com. (2017). *Flexify 2 Photoshop plug-in*. [online] Available at: http://www.flamingpear.com/flexify-2.html [Accessed 20 May 2017].
fxguide. (2017). *Tandent Lightbrush: engineering an image*. [online] Available at: https://www.fxguide.com/featured/tandent-lightbrush-engineering-an-image/ [Accessed 27 Apr. 2017].
fxphd. (2017). *fxphd*. [online] Available at: https://www.fxphd.com [Accessed 11 Apr. 2017].
Gdcvault.com. (2017). *Photogrammetry and 'Star Wars Battlefront'*. [online] Available at: http://www.gdcvault.com/play/1023272/Photogrammetry-and-Star-Wars-Battlefront [Accessed 10 Feb. 2017].
Geomagic.com. (2017). *Geomagic Design X*. [online] Available at: http://www.geomagic.com/en/products-landing-pages/designx [Accessed 14 May 2017].
GIS Lounge. (2017). *A Brief Introduction to Photogrammetry and Remote Sensing ~ GIS Lounge*. [online] Available at: https://www.gislounge.com/a-brief-introduction-to-photogrammetry-and-remote-sensing/ [Accessed 9 Aug. 2017].
Guillermoluijk.com. (2017). *DCRAW TUTORIAL*. [online] Available at: http://www.guillermoluijk.com/tutorial/dcraw/index_en.htm [Accessed 30 Jan. 2017].
Hart, F. (2017). *Smart Shooter Photography Software*. [online] Kuvacode.com. Available at: https://kuvacode.com/ [Accessed 26 Apr. 2017].
Hoyafilterusa.com. (2016). *HOYA SOLAS IRND*. [online] Available at: https://hoyafilterusa.com/product/hoya-solas-irnd-2/ [Accessed 28 Oct. 2016].
Igus.eu. (2017). *igus drylin*. [online] Available at: http://www.igus.eu/drylin [Accessed 11 Aug. 2016].
Illumni.co. (2017). *Using the Unity Game Engine to Design Interactive Lighting Simulations*. [online] Available at: http://www.illumni.co/using-the-unity-game-engine-to-design-interactive-lighting-simulations-case-kruunuvuorensilta-bridge-helsinki-finland [Accessed 2 Jul. 2017].
Inanici, M. (2017). *Mehlika Inanici - Research*. [online] Faculty.washington.edu. Available at: http://faculty.washington.edu/inanici/MI-RESEARCH.html [Accessed 13 Oct. 2016].
inc., L. (2017). *3D Laser Scanner, Coordinate Measuring Machine and 3D Scanning | Creaform*. [online] Creaform3d.com. Available at: https://www.creaform3d.com/en [Accessed 15 Jul. 2017].
Ipb.uni-bonn.de. (2017). *Photogrammetry I & II – IPB*. [online] Available at: http://www.ipb.uni-bonn.de/photogrammetry-i-ii/ [Accessed 10 May 2017].
Kuka.com. (2017). *KUKA AG*. [online] Available at: https://www.kuka.com/ [Accessed 15 Jul. 2016].
Kumaragurubaran, V. (2013). *High Dynamic Range Image Processing Toolkit for Lighting Simulations and Analysis*. 1st ed. [ebook] University Of Washington, pp.1-71. Available at: http://courses.washington.edu/hdrscope/documentation/viswa_thesis.pdf [Accessed 8 Sep. 2016].
Marmoset. (2016). *HDR Panorama Photography | Marmoset*. [online] Available at: https://www.marmoset.co/posts/hdr-panorama-photography/ [Accessed 23 Jun. 2017].
McHugh, S. (2016). *Understanding photography*. 1st ed. Cambridge in Colour.

NVIDIA Developer. (2016). *GPU Gems*. [online] Available at: https://developer.nvidia.com/gpugems/GPUGems3/gpugems3_ch24.html [Accessed 30 Oct. 2016].

Orientalmotor.com. (2017). *Stepper Motors*. [online] Available at: http://www.orientalmotor.com/stepper-motors/mechanical-absolute-encoder-az-series.html [Accessed 6 Jan. 2017].

Peterson, B. (2004). *Understanding exposure*. New York: Amphoto Books.

Pharr, M. and Humphreys, G. (2004). *Physically based rendering*. 1st ed. Amsterdam [u.a.]: Elsevier [u.a.].

Pixologic ZBrush. (2017). *Pixologic : ZBrush*. [online] Available at: http://pixologic.com/ [Accessed 4 Feb. 2017].

Ptgui.com. (2017). *PTGui Stitching Software*. [online] Available at: https://www.ptgui.com [Accessed 28 Jan. 2017].

Radiance-online.org. (2017). *Radiance - A Validated Lighting Simulation Tool — Radsite*. [online] Available at: https://www.radiance-online.org/ [Accessed 9 Nov. 2016].

Salvaggio, N., Stroebel, L. and Zakia, R. (2013). *Basic photographic materials and processes*. 3rd ed.

Sébastien Lagarde. (2015). *Siggraph 2014 : Moving Frostbite to Physically based rendering V3*. [online] Available at: https://seblagarde.wordpress.com/2015/07/14/siggraph-2014-moving-frostbite-to-physically-based-rendering/ [Accessed 11 Mar. 2017].

Sekonic.com. (2016). *Sekonic Light Meter*. [online] Available at: http://www.sekonic.com/denmark/products/l-478d/overview.aspx [Accessed 3 Sep. 2016].

Sharonov, O. (2016). *3D LUT Creator*. 3D LUT Creator.

Stereolabs.com. (2017). *ZED Stereo Camera*. [online] Available at: https://www.stereolabs.com/ [Accessed 27 Jul. 2017].

Stockist, A. (2017). *Aluminium - Specifications, Properties, Classifications and Classes, Supplier Data by Aalco*. [online] AZoM.com. Available at: http://www.azom.com/article.aspx?ArticleID=2863 [Accessed 15 Oct. 2016].

Support.solidangle.com. (2016). *Arnold for Maya*. [online] Available at: https://support.solidangle.com/display/AFMUG/Render+Selection+to+Texture [Accessed 21 May 2017].

Tinkerforge.com. (2017). *Home | Tinkerforge*. [online] Available at: https://www.tinkerforge.com/en/ [Accessed 3 Jul. 2016].

Unity. (2017). *Unity - Game Engine*. [online] Available at: https://unity3d.com/ [Accessed 30 Jan. 2017].

Unrealengine.com. (2017). *Unreal Engine 4*. [online] Available at: https://www.unrealengine.com/what-is-unreal-engine-4 [Accessed 11 Jan. 2017].

Verhoeven, G. (2016). *Basics of photography for cultural heritage imaging*. [online] Hdl.handle.net. Available at: https://biblio.ugent.be/publication/8050621 [Accessed 14 Nov. 2016].

Xnormal.net. (2017). *xNormal web page*. [online] Available at: http://www.xnormal.net/ [Accessed 8 May 2017].

Xritephoto.com. (2017). *ColorChecker Passport*. [online] Available at: http://xritephoto.com/colorchecker-passport-photo [Accessed 30 Aug. 2016].

YouTube. (2017). *NVIDIA's Reality Capture / Unity Photogrammetry Viewer*. [online] Available at: https://www.youtube.com/watch?v=6SD3hZgQpnY [Accessed 27 Jul. 2017].