Assignment 1

Photogrammetry

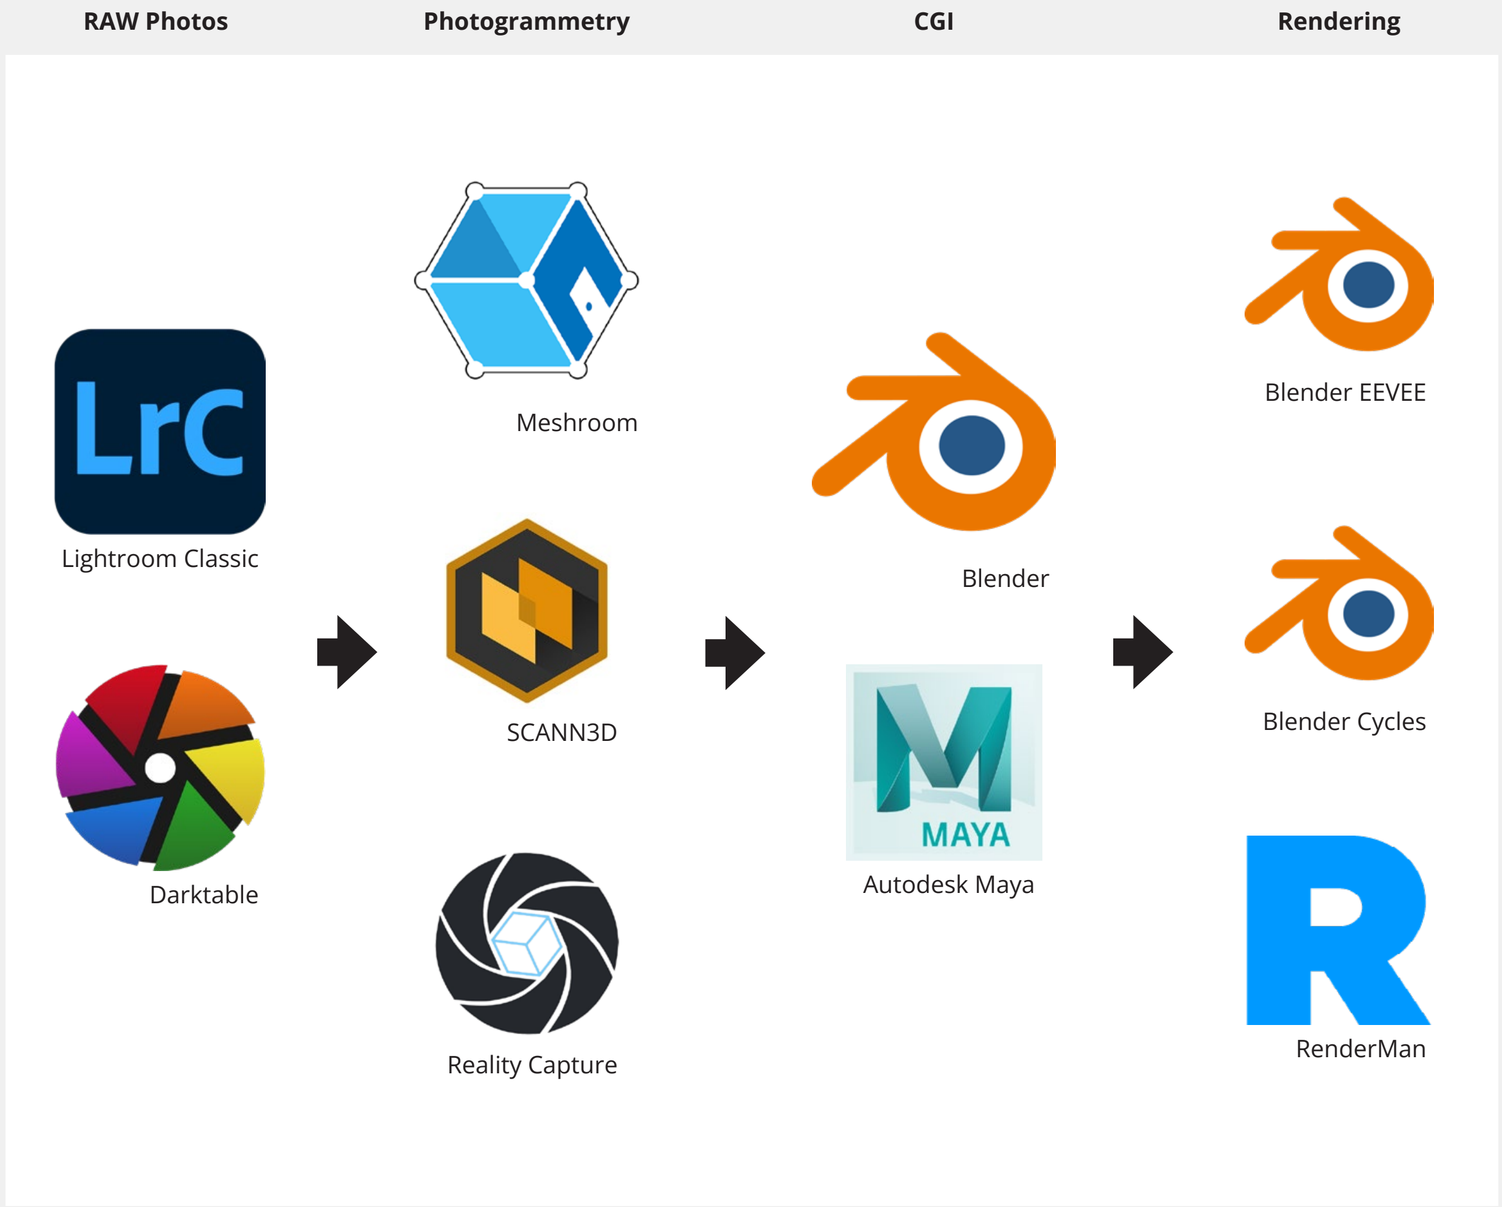

Assignment 1: Photogrammetry

Software and Process Overview

In order to make a successful photogrammetry object, I needed a variety of programs.

The process starts with using a DSLR camera to capture RAW photos with consistent exposures. These are brought into a photo developer application such as Adobe Lightroom or the free open-source Darktable. These are adjusted to better show detail, before being exported to a more appropriate format, such as JPG.

The developed photos are imported into a photogrammetry program, such as RealityCapture, which creates a point cloud from the batch of images. This cloud is then used to create the object's geometry and textures.

Once done, I exported the OBJ geometry and PNG texture files. My justification for these specific formats is that OBJ is a relatively minimal option for keeping the geometry, UVs and not much else. PNG is a versatile format that remains fairly small without losing quality.

Software Screenshots

SCANN3D user interface

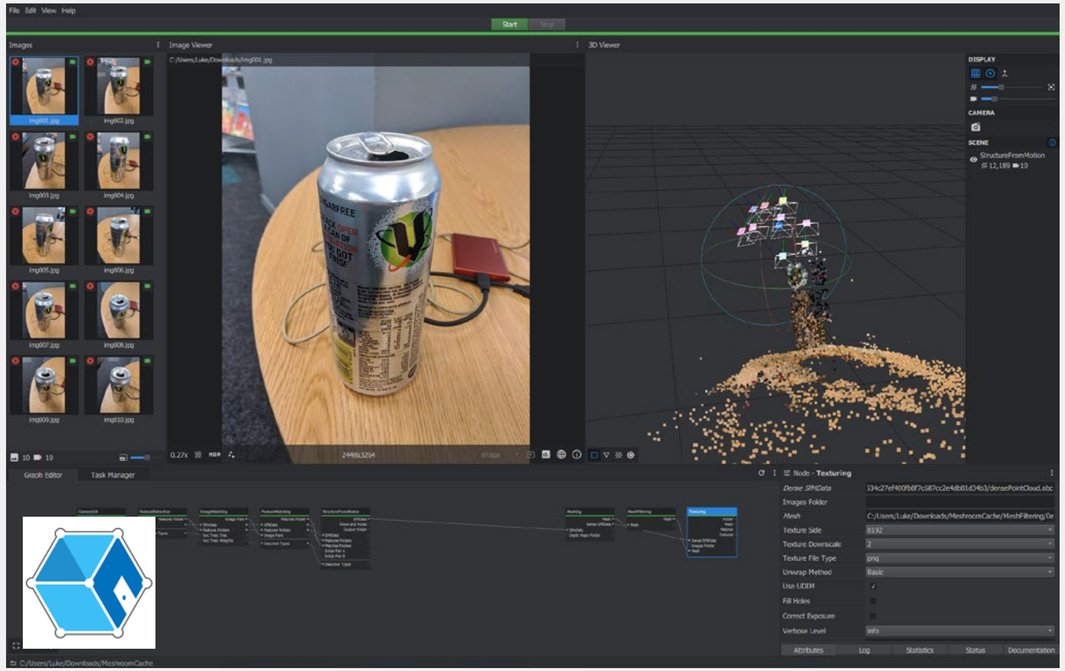

Meshroom user interface

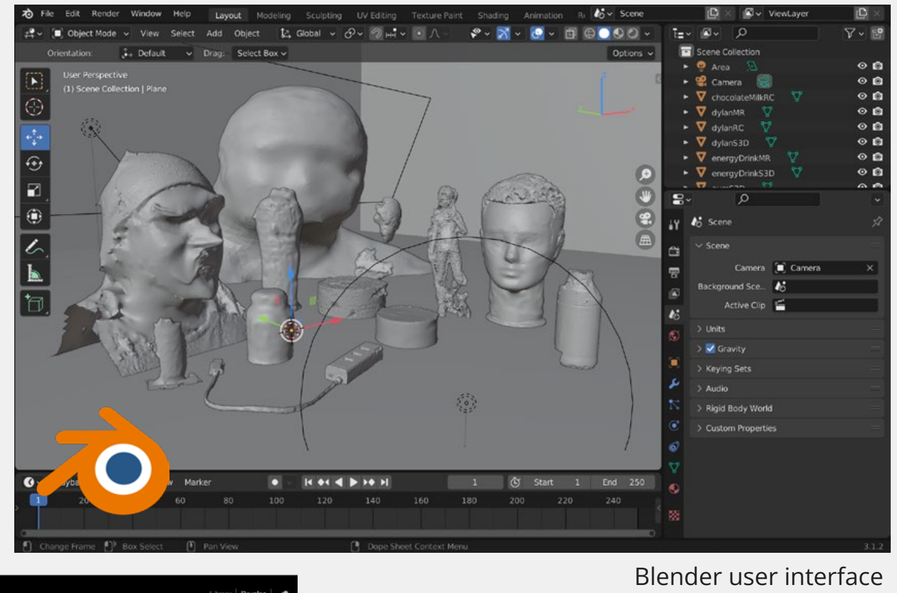

Blender user interface

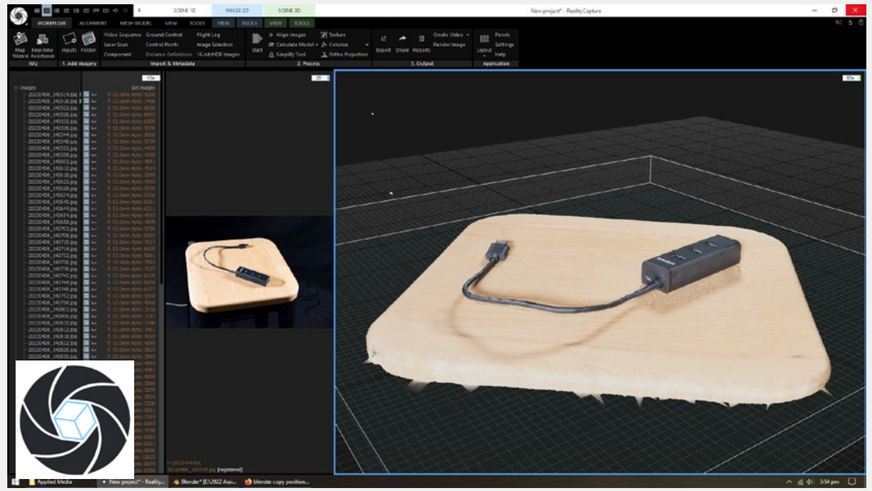

RealityCapture user interface

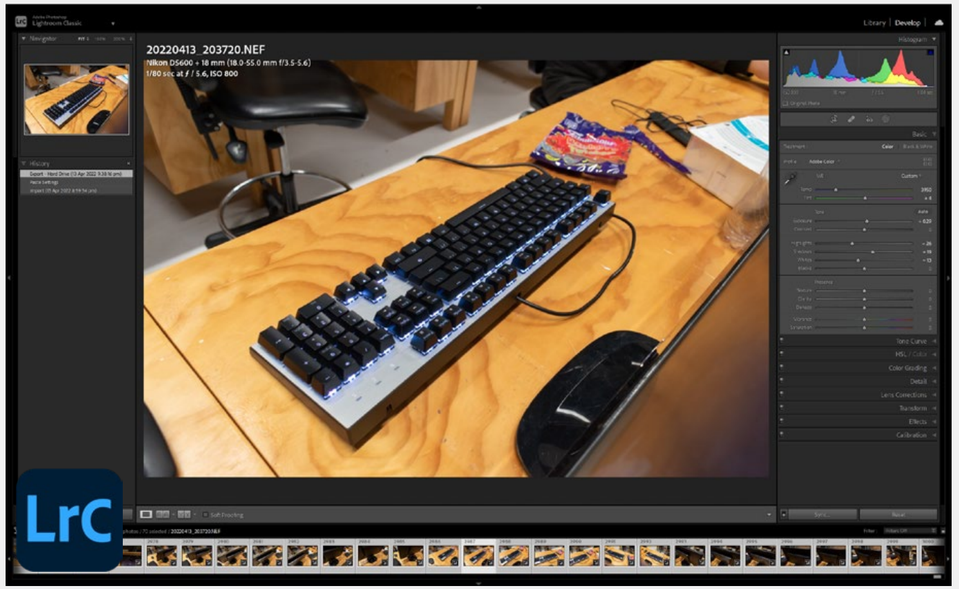

Adobe Lightroom user interface

Texture Maps

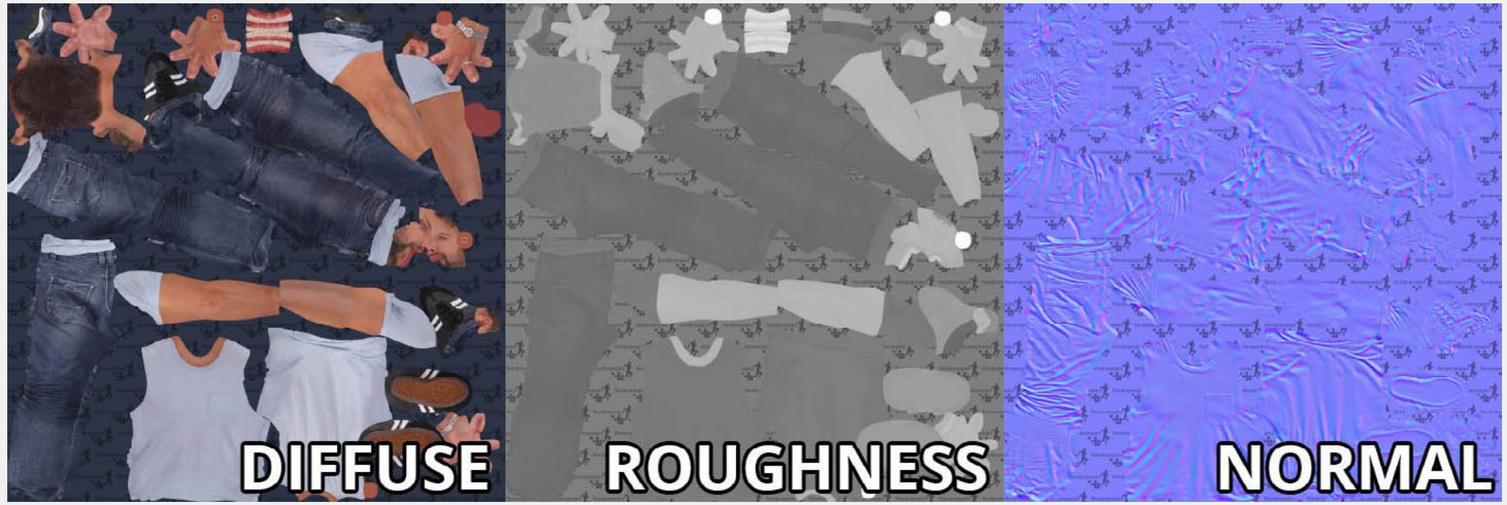

Textures sourced from: https://free3d.com/3d-model/nathan-animated-003-walking-644277.html

Diffuse:
Standard color map with shadows. This is most likely what you'll get from photogrammetry software.

Albedo:
Also a color map but with no lighting or shadows. The result is a very flat looking texture.

Ambient occlusion:
Typically a greyscale texture that identifies lighting and shadows. This placed over an albedo map equals a diffuse map.

Bump/normal:
Used to fake detail without changing the geometry itself. Bump maps are greyscale, while normal maps are generally blueish.

Displacement:
Changes the geometry of the model. More "expensive" than bump/normal maps but produces more realistic results.

Roughness:
Typically greyscale. Controls the roughness/glossiness of the material.

UV layout:
Not a texture per se, but more of a guide for how texture maps translate to the mesh.

A folder of rendered image sequences saved in the EXR file format

Best Practices

When it comes to digital media, I believe it's useful to have a good grasp on the various formats and standards.

Source images:
The photos taken with my DSLR are shot in RAW with consistent exposure settings. Once I take these into Lightroom, I remove the natural lens distortion and vignetting and export the photos as JPGs.

Texture maps:
As good practice, I keep my textures saved as PNGs, ensuring they won't degrade in quality after each save (as is the case with JPGs) and with a bit depth of 16 per channel, to reduce the risk of color banding.

Objects:
For cross-compatibility between Maya, Blender and other 3D programs, I decide to save my objects as OBJ files. This small and versatile format contains geometry data, UVs and not much else.

Rendered images:
When rendering from Maya/Blender, I make sure to use the 16-bit EXR format to preserve information in case colour correction is needed. Think of this as the CGI equivalent to RAW photos from a DSLR.

Lightroom Tricks

On the day the assignment was due, I figured it made sense to try compressing the highlights and shadows of the source images in Lightroom, both to retain the detail in the extremes when exporting the JPGs, and to minimize the effects of the lighting as much as possible. Unfortunately, this particular set of images shown to the left didn't work at all.

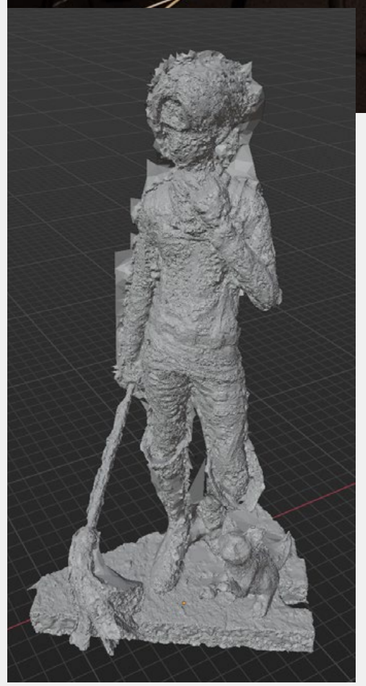

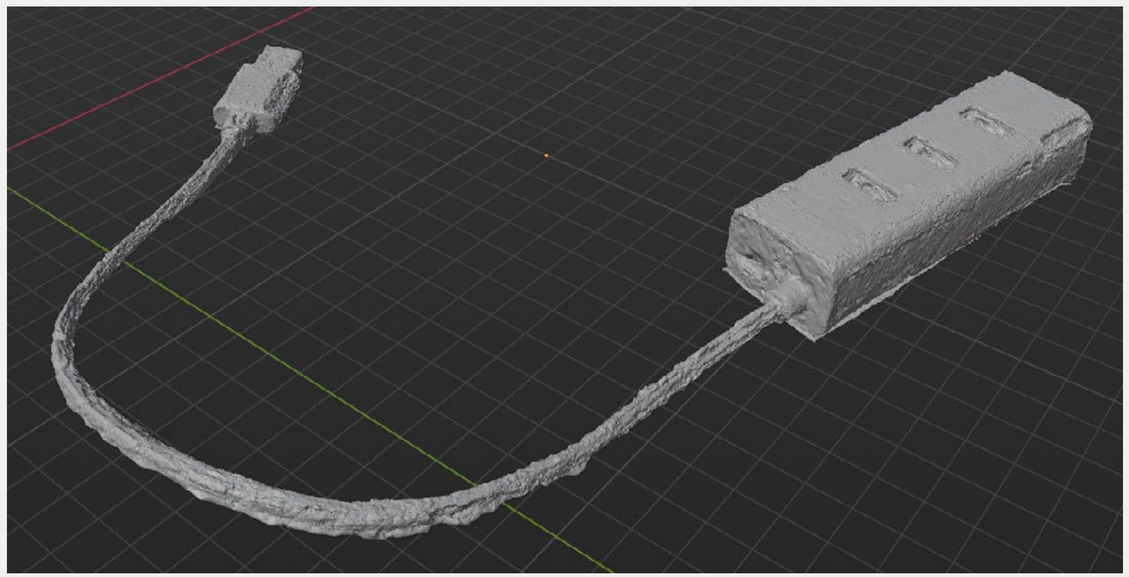

Assignment 1: Photogrammetry

Controlled Environment Tests

I'm no expert in lighting, but I tried making 2 models with studio floodlights and a thick black curtain in the background. The setup was slightly different for each, with reflective sheets used for the janitor object.

USB Hub
Camera: Nikon D5600
Exposure: 1/80s f/4.5 ISO 800

The lighting setup for this was similar to the photo shown, only without the reflective sheets or a tripod. Despite the object being fairly small, there were still hard shadows present. For the shots, I moved around the object in 3 rotations, with a few random close-ups for the sake of providing detail.

Janitor
Camera: Nikon D7200
Exposure: 1/50s f/5.3 ISO 1600

For this test, the camera was mounted on a stationary tripod, with the object itself being rotated by another student on command. The tripod's height was adjusted after a full rotation.

Natural Lighting Tests

Camera: Nikon D5600
Exposure: 1/100s f/5.0 ISO 800

This test of Dylan was deliberately taken in an unevenly-lit environment. This was the only object to be created in all 3 photogrammetry programs (SCANN3D, Meshroom and RealityCapture). RealityCapture managed to recreate most of his face, but Meshroom struggled with the side shown, which I assume was due to the overexposed background.

Dylan objectified by SCANN3D, Meshroom and RealityCapture

Assignment 1: Photogrammetry

Phone Camera Test

Camera: Google Pixel 4
Exposure: Inconsistent

I did a quick test of a chocolate milk bottle using my phone camera, which had the added convenience of faster autofocus amongst other things.

Assignment 1: Photogrammetry

Environment Test

Camera: Nikon D5600
Exposure: 1/80s f/5.6 ISO 800

I took some photos from a section of the classroom. Although it was evening, the lighting was good, and I just clicked away from random angles.

The resulting object (made in RealityCapture), was tens of millions of polygons. Even the computers at school equipped with 32GB of RAM only just managed to import it.

I decimated the model (as shown before), then deleted all faces except for what's shown to the left.

I figured this would be a good background to use for stuff.

This particular test also made me realize how convenient photogrammetry could be for properly remodelling a real-life scene. It eliminates the need for tape measures and getting proportions right. The basin shown for instance could be optimized with simple planes using what's already there as reference.

DSLR Exposure Settings

Stuff that feels like commonsense to a long-time photographer like me. Fast shutter speeds when holding the camera result in sharper shots, which are important for the software to pick up details. Depending on various things like wind and caffeine jitters, I found 1/80s to be the slowest ideal choice.

ISO could go up to around 800 before noise crept in. Not a huge deal as Lightroom's denoiser works well.

For the aperture, I tried to leave it between f/5.6-9.0 where possible, to get enough depth in focus without crossing the line into lens diffraction at higher f-stops.

Photos were shot in RAW as always, but after developing were exported as JPGs because that's all that was needed.

DSLRs are generally a better choice over phone cameras, although I found the focus-locking on my Pixel 4 useful when shooting the chocolate milk.

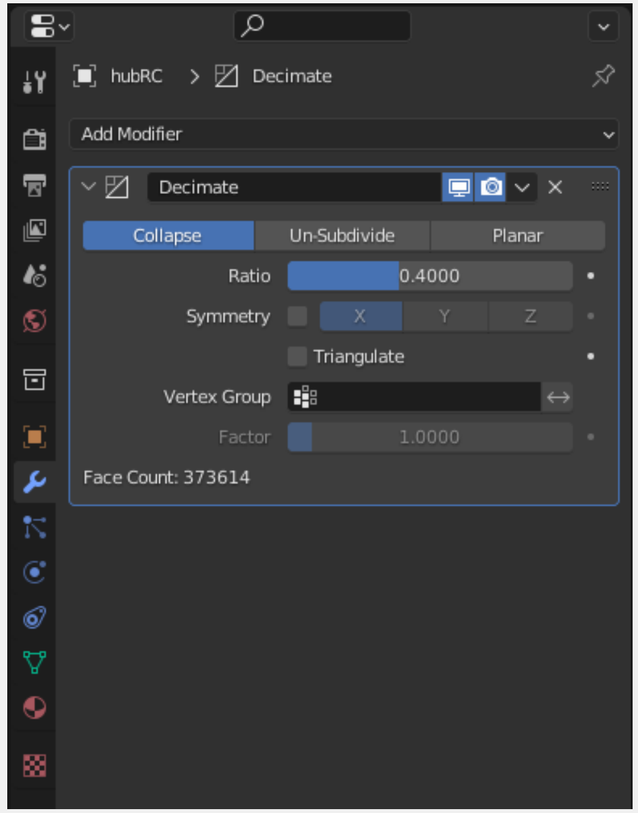

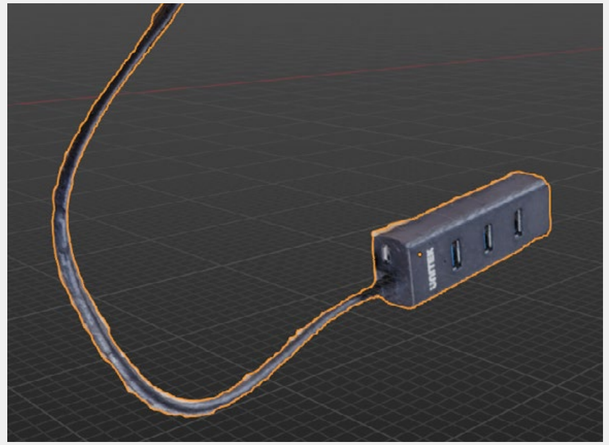

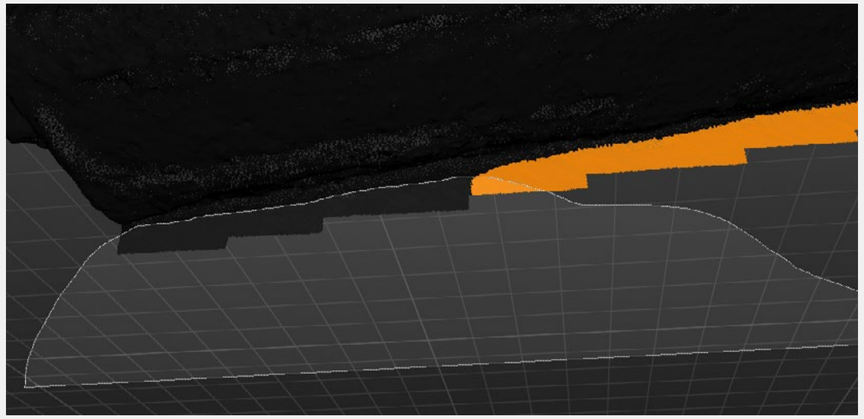

Assignment 1: Photogrammetry

Blender Cleanup

Decimate
The first step after bringing the objects into Blender was to reduce the polycount. Some of the objects were very dense, namely Dylan's head, which had 12 million faces. This tool kept the overall shape without affecting the textures. The slider of interest, labelled Ratio, considers a value of 1 to be the original density, ergo I found a value of 0.3 to be adequate for most models.

Polygon deletion
The photogrammetry software would often consider bits of the background to be part of the object, and would include them in the model. To fix this, I simply switched to Edit Mode and deleted the unnecessary faces from around the subject.

Reshaping/Smoothing
By using Blender's sculpting tools, I was able to smooth the surface of the models to reduce the bumpy bits. This in combination with the Smooth Shading feature resulted in a cleaner model.

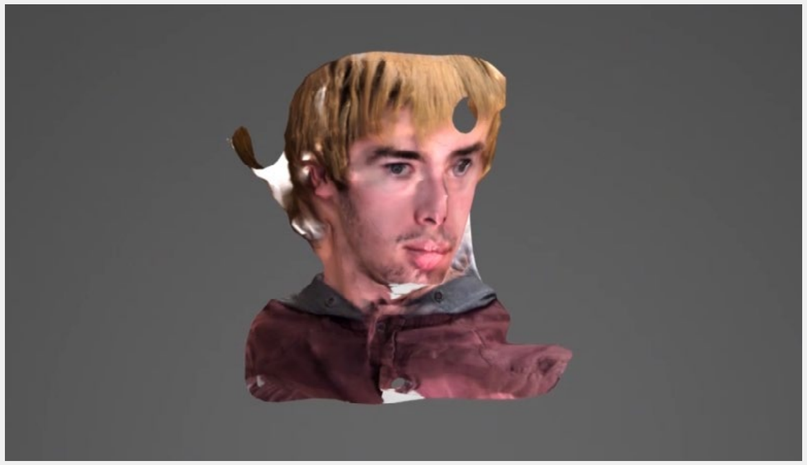

Photogrammetry test back in 2017

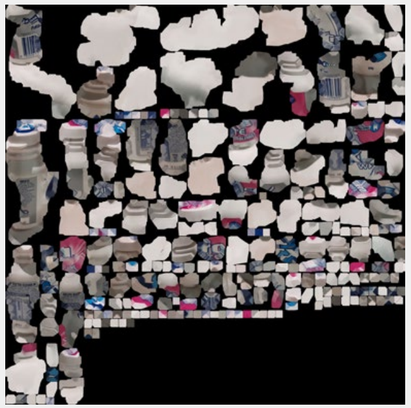

Gum texture

Gum preview

Source images

Gum shown in the Unity WebGL player

Assignment 1: Photogrammetry

2017 and 2019 Tests

Turns out I did photogrammetry twice before, once in 2017 when I didn't even know what it was called, then again in 2019 for an assignment.

All I have of the former is a rendered video, but from what I remember this only used 5 shots from my phone's selfie camera, which would explain the quality. I had brought this into Maya and done a simple gravity simulation.

The gum was made using SCANN3D, which would then be brought into a simple Unity model viewer of my own making and exported for the WebGL platform.

Final Renders

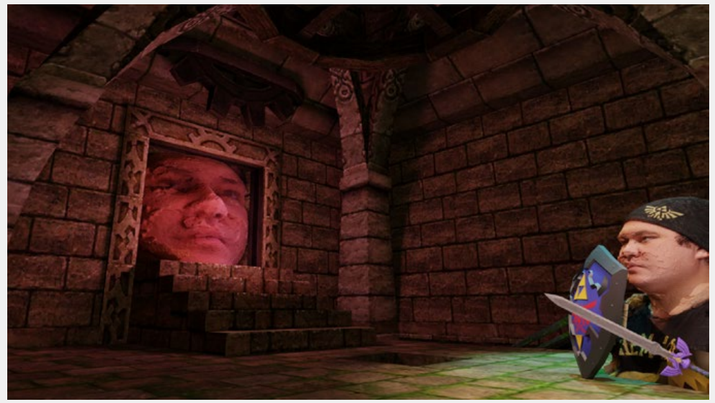

Assignment 2

CGI Character and Environment

Rationale and overview

Following from the success of his head from assignment 1, I asked Dylan if he wanted to participate in a full-body shoot.

The shoot took 10 minutes, and as a result Dylan's arms lowered from fatigue. This caused RealityCapture to entirely omit them. To fix this, I simply did another shoot next to a table.

One of the main appeals of photogrammetry for me is the ability to use a rough generated model as a base for retopology, just like how a digital sculptor would when using ZBrush for instance. This eliminates the need for front and side reference images.

Having Dylan recreated in CGI form would be useful for the short film or render I have planned for next semester.

Assignment 2: CGI Character and Environment

Generating the reference model

Like with assignment 1, I took many photos around the subject covering all areas. I asked Dylan to look his best, and we used the area behind campus as it seemed like a good place to get even lighting.

The model was generated in RealityCapture with 107 high-quality images.

Specs:
- Nikon D5600 with 18-55mm lens
- 1/100s shutter
- f/5.6 aperture
- ISO 400

The images were taken around 10:30 in the morning on the 11th of May.

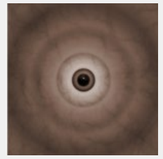

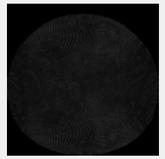

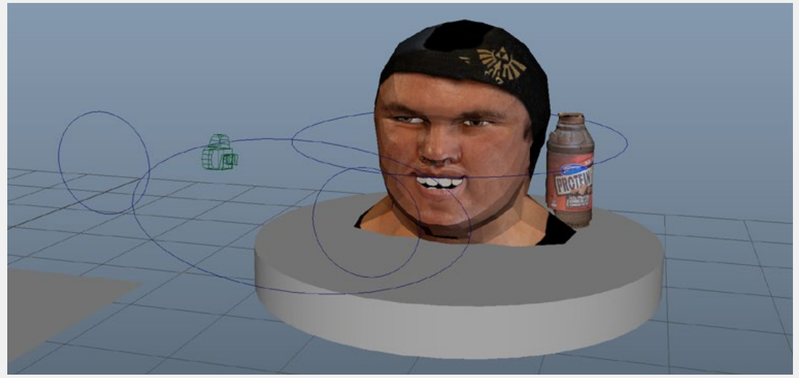

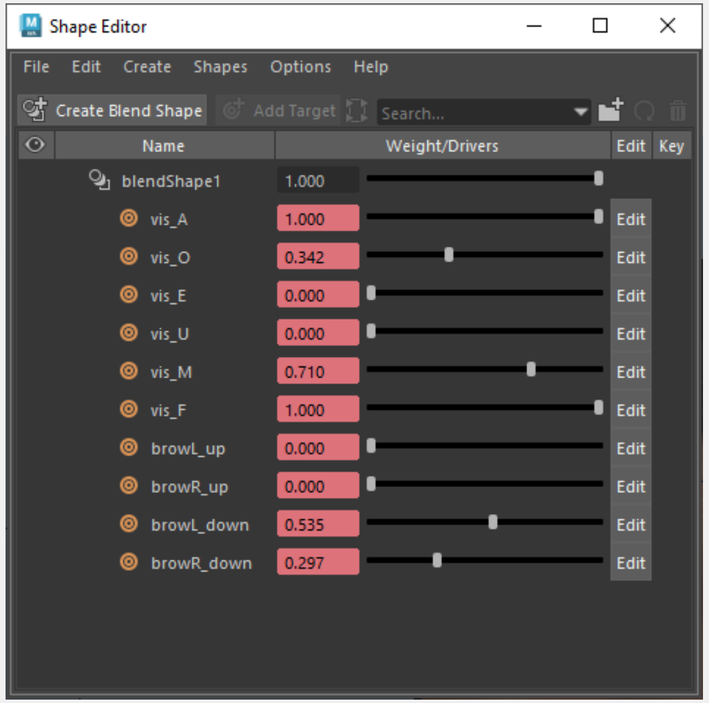

Assignment 2: CGI Character and Environment

Uncanny animation prank

To test my facial animation skills, I created clean topology of Dylan's head and applied an edited texture from the RealityCapture model.

The only form of rigging used were NURBS circles to rotate the head and point the eyes. All other movement was used with blendshapes.

The result is a very disturbing short animation. Viewer discretion is advised.

Retopology

I used Maya's Quad Draw tool to create sensible humanoid topology based on Dylan's shape, so that the character could be easily deformed and textured. Some areas were rushed for time, namely the elbows, knees and neck.

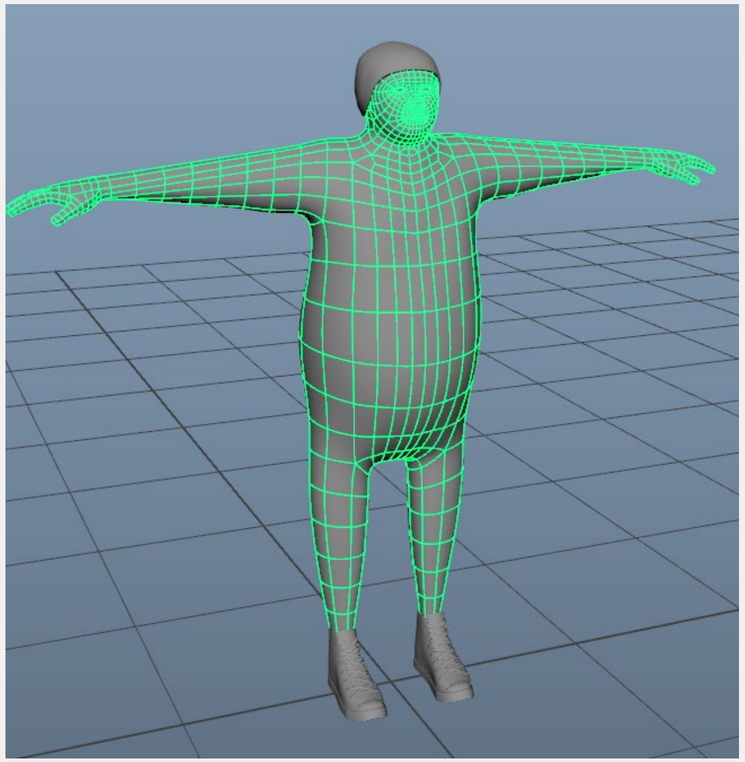

Texturing

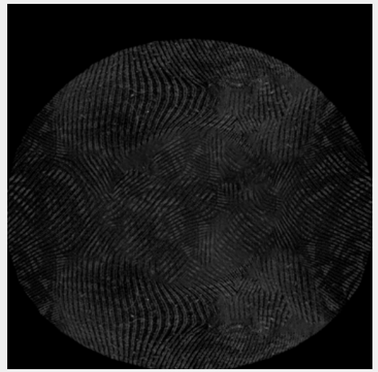

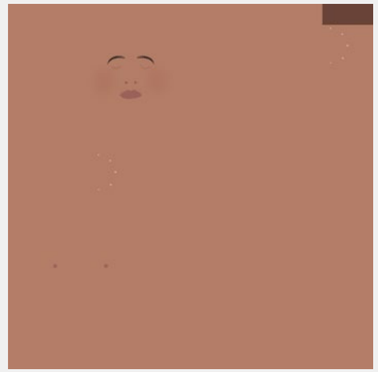

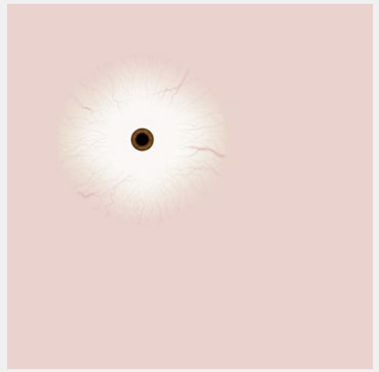

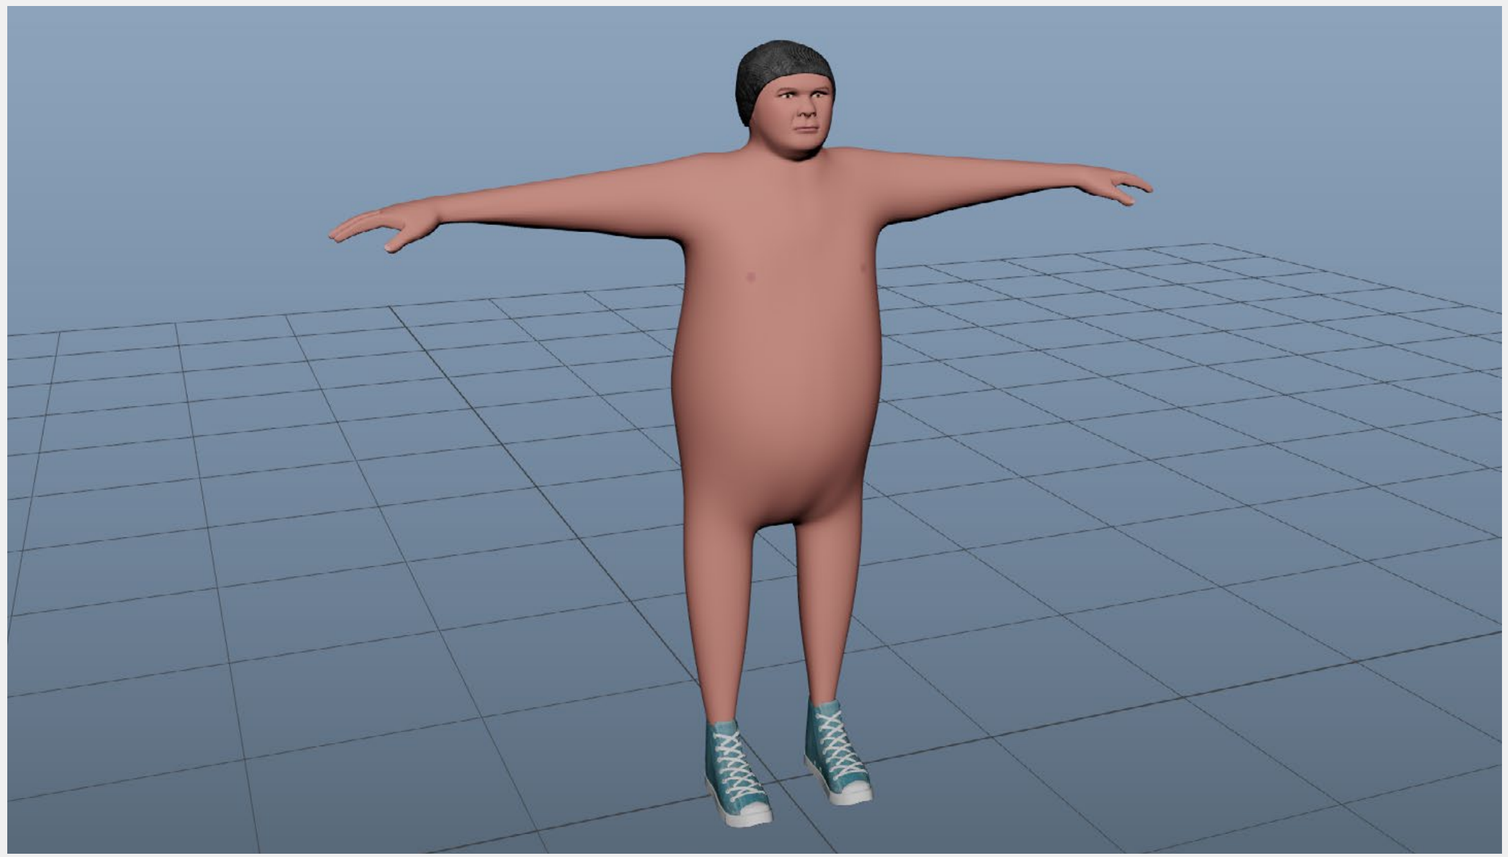

For the sake of this assignment, I kept the textures simple, creating only basic diffuse maps for the beanie, body, eyes and shoes.

The body texture features eyebrows, eyeliner, cheeks, nostrils, lips, fingernails and nipples.

Cloth simulation

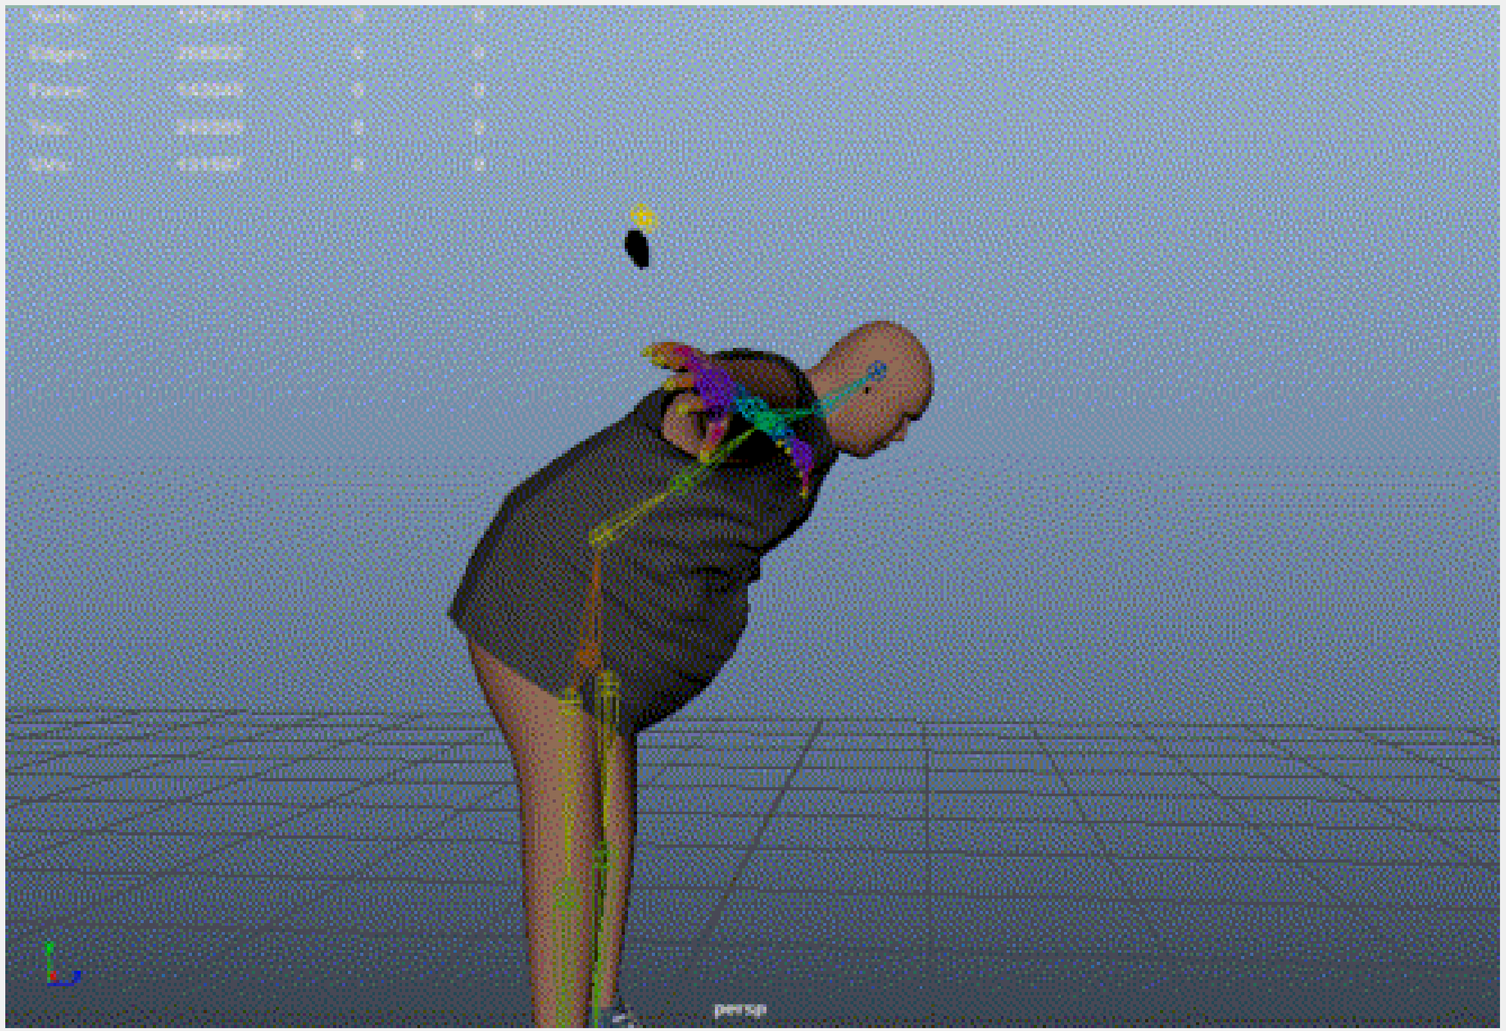

I experimented with simulating fabric using nCloth, a feature in Autodesk Maya.

Not much else to say here.

Final render

Here's the final render of Dylan in a simple showcase environment. Compared to Uncanny Dylan, this version is much less frightening.

Assignment 2: CGI Character and Environment

Rationale and overview

This assignment required a full 3D environment, and while it could be something simple, I decided to really put in the effort to recreate the scene pictured to the left, located at Mount Victoria, roughly a 5 minute walk south of the summit.

Before getting started, I was discouraged by the complexity of the trees and bushes in the background. How would I recreate all of them?

Visual moodboard

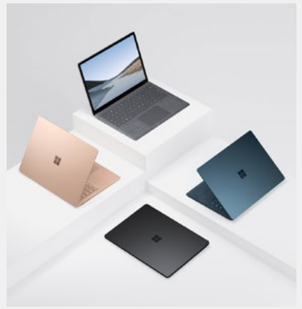

Here are two images created by Microsoft.

The top one is reminiscent of Windows XP's well-known Bliss wallpaper, albeit with softer lighting and colours. The level of detail left me unsure of whether it was a render or a real photo, but regardless, the style is the reason I included it.

The bottom one was released alongside Microsoft's Surface Laptop 3 (a product I do not recommend). This stands out more as a render due to the unrealistic nature of it. Interesting to note is the colour of the grass, which aligns with the colours the laptop was released in.

Together, both images share characteristics that are soft, calm and simple.

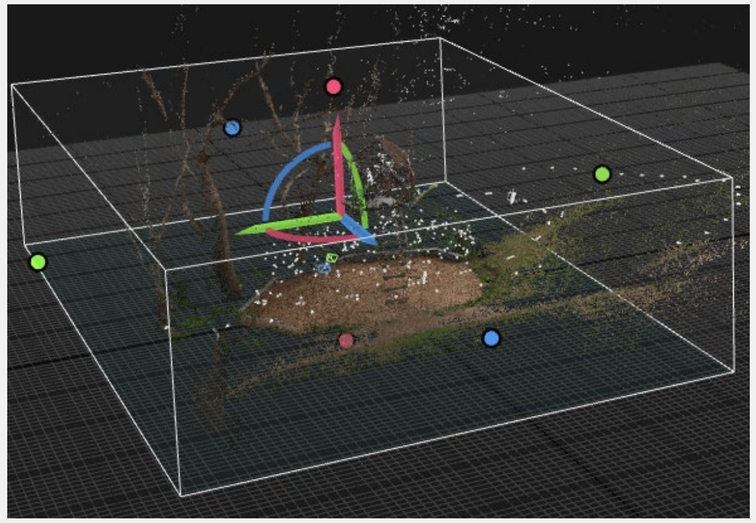

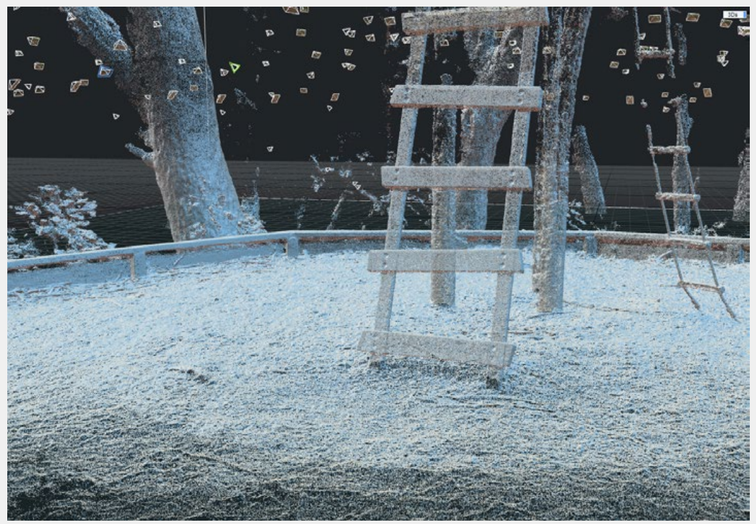

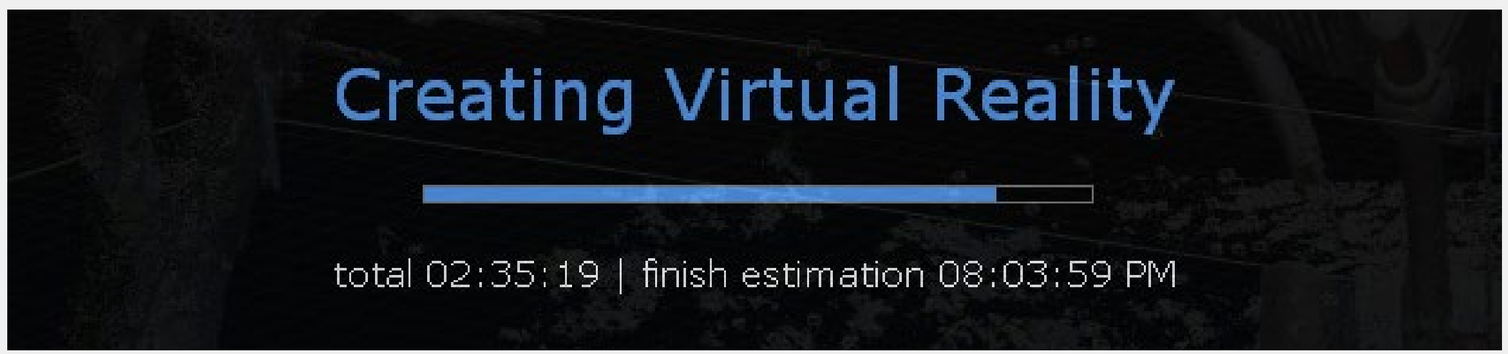

Assignment 2: CGI Character and Environment

Generating the reference model

Working with such a large scene required capturing every inch of the location. Over a thousand raw images were taken, but not before maxing out my camera's 32GB card, resulting in me having to return another day.

Bringing the photos into RealityCapture worked, but it took a very long time to process, and required simplifying the model before I could even export it!

Remodelling

Like I demonstrated in an earlier page, this model required clean topology before it could pass my standards. Considering the ladder pictured, you can see just how much lower the polycount is, reducing the strain on my computer while fixing the model's shape and structure.

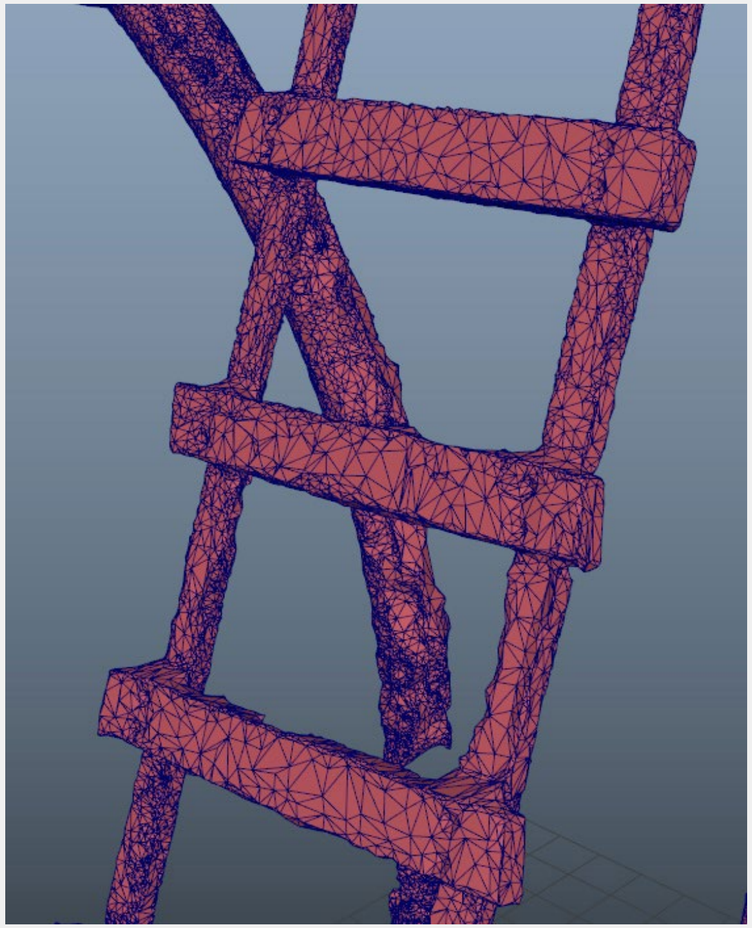

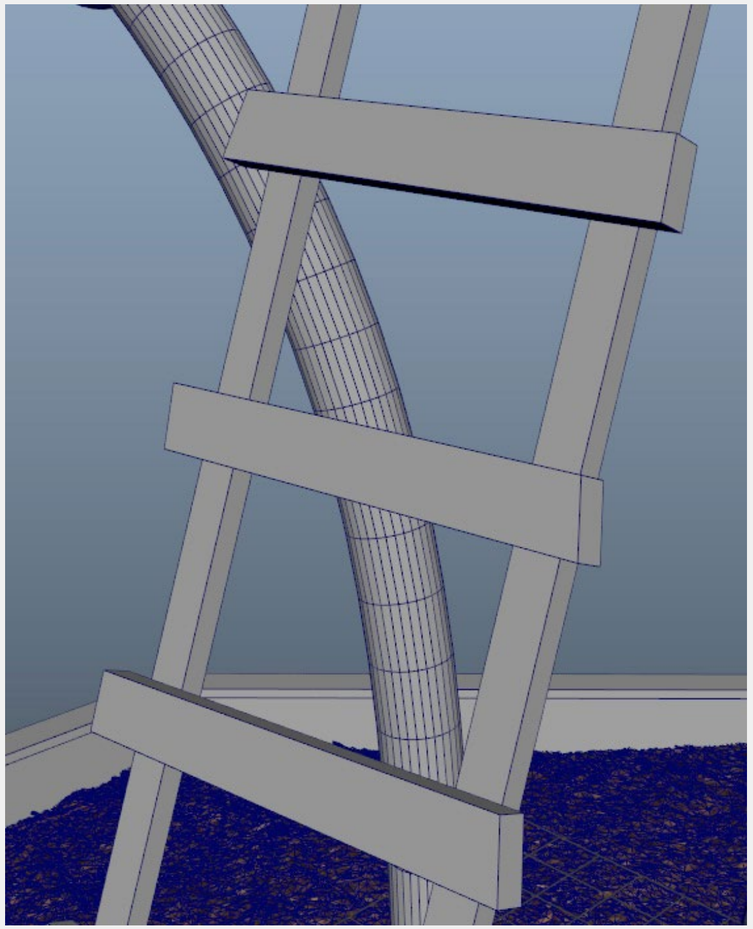

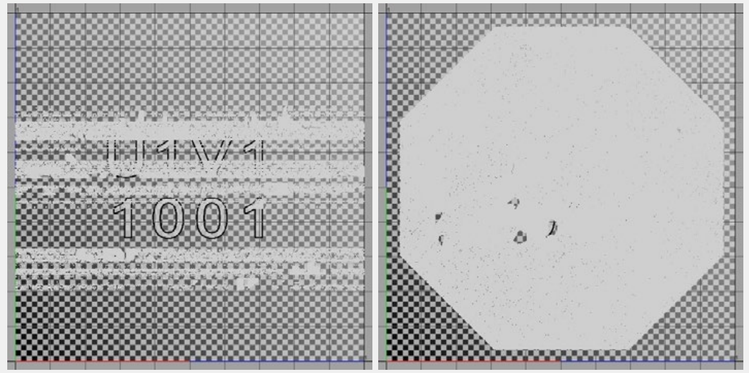

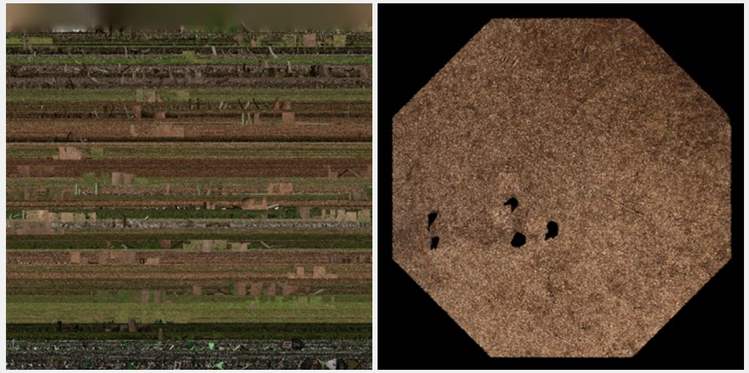

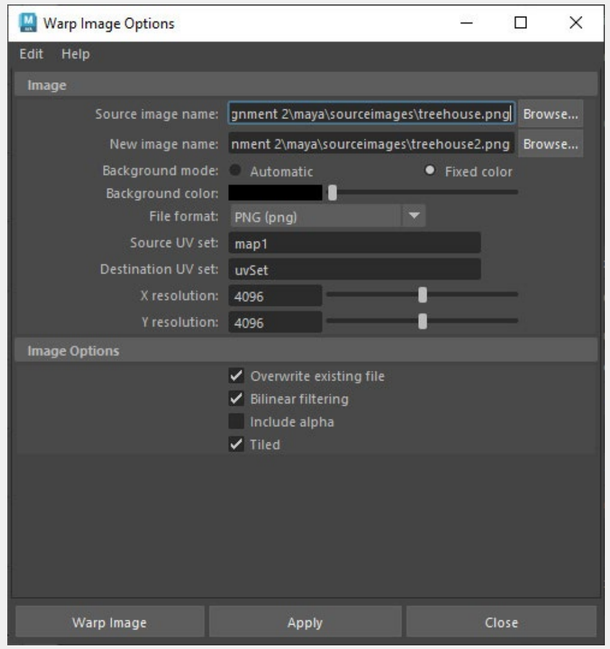

Assignment 2: CGI Character and Environment

Texturing tricks

I found a neat feature hidden in Maya's UV Editor window called Warp Image. The gist of how this works is you create a second UV Set with different layouts to the first (top-left screenshots). Maya uses this information to rearrange the source image data (top-right screenshots). I used this specifically to create a tidy bark texture.

I also edited the texture for the ladder's steps to create specular and displacement maps for the wood polish and studs respectively. Small things to add detail.

Grass: PaintFX or XGen?

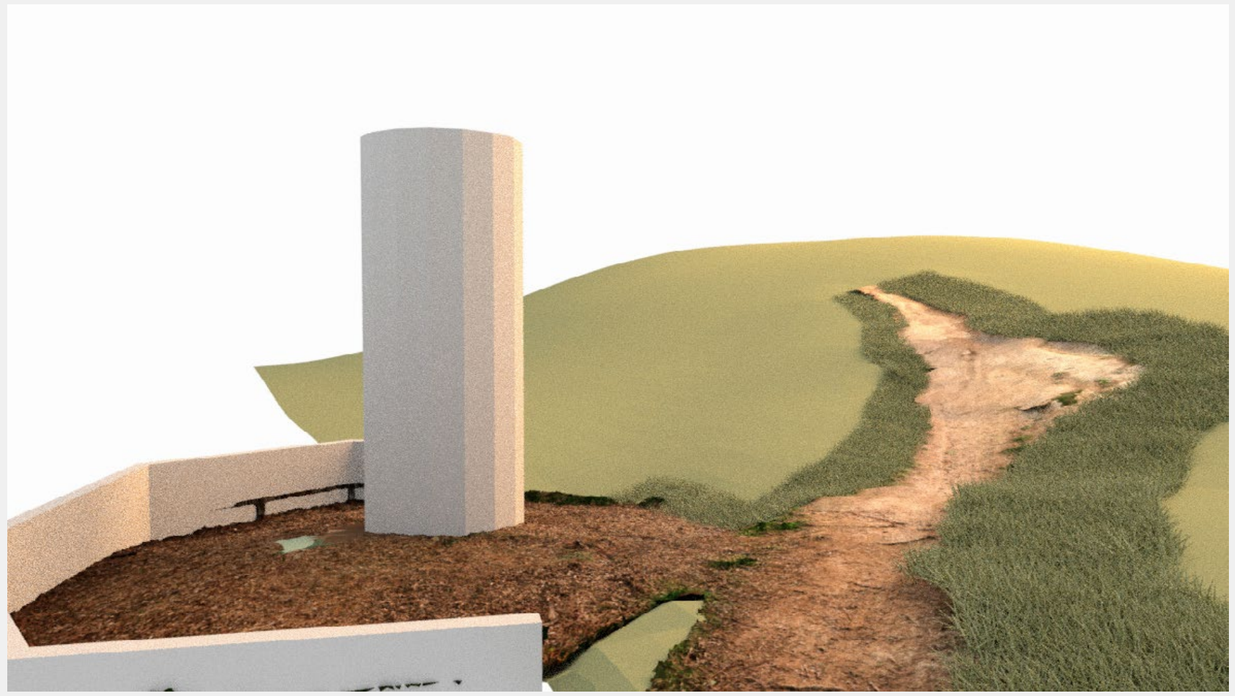

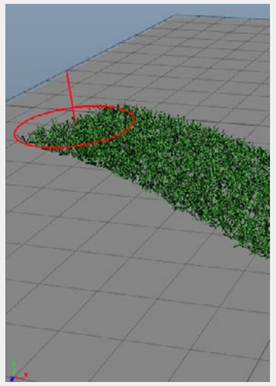

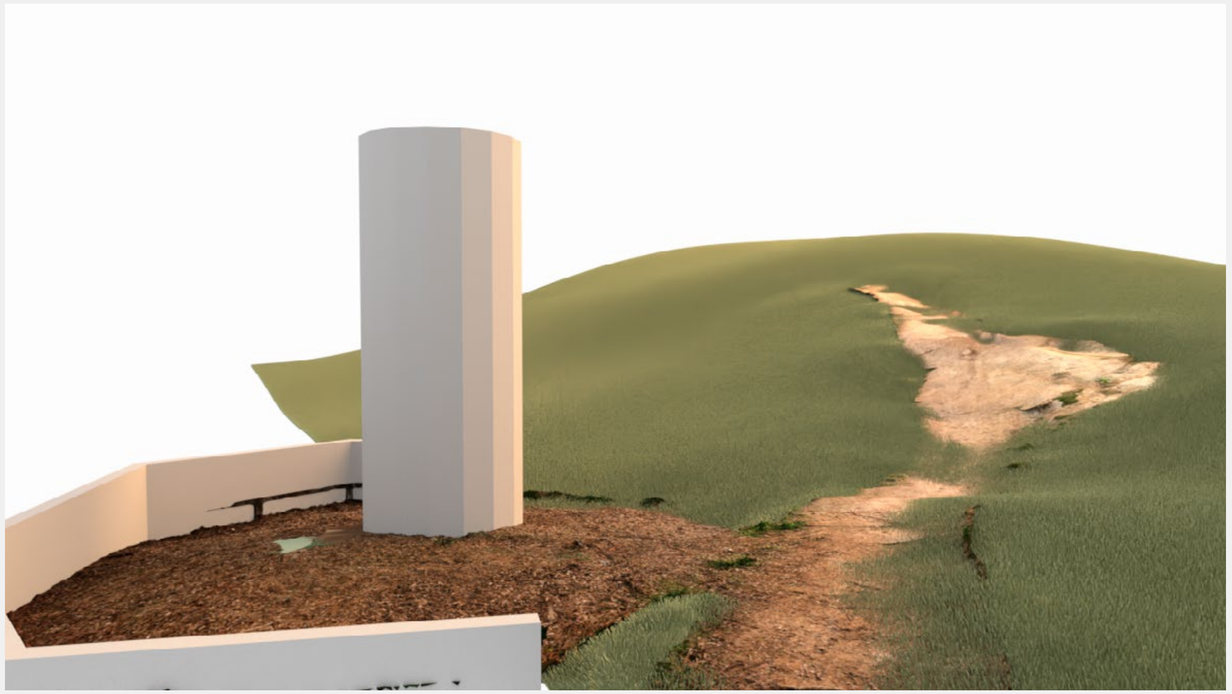

When it came time to create the grass, I had 2 methods at my disposal which work different ways.

The first method, PaintFX, involves choosing a pre-built mesh and painting it on the surface of the hills. Various attributes of the grass could be changed, but my understanding is this method leaves the grass stiff and hard to animate, if not impossible.

The second method, XGen, involves dynamically generating hair, fur or in this case grass, by controlling various attributes through ramps, expressions and texture maps. It can also be connected to a nHair system to make it respond to wind.

I decided to go with XGen. I used an edited version of the hill texture and assigned it to the density and length attributes, resulting in the dirt path being free of grass. The white areas of the texture indicate presence, while the black areas indicate absence.

Posing Dylan

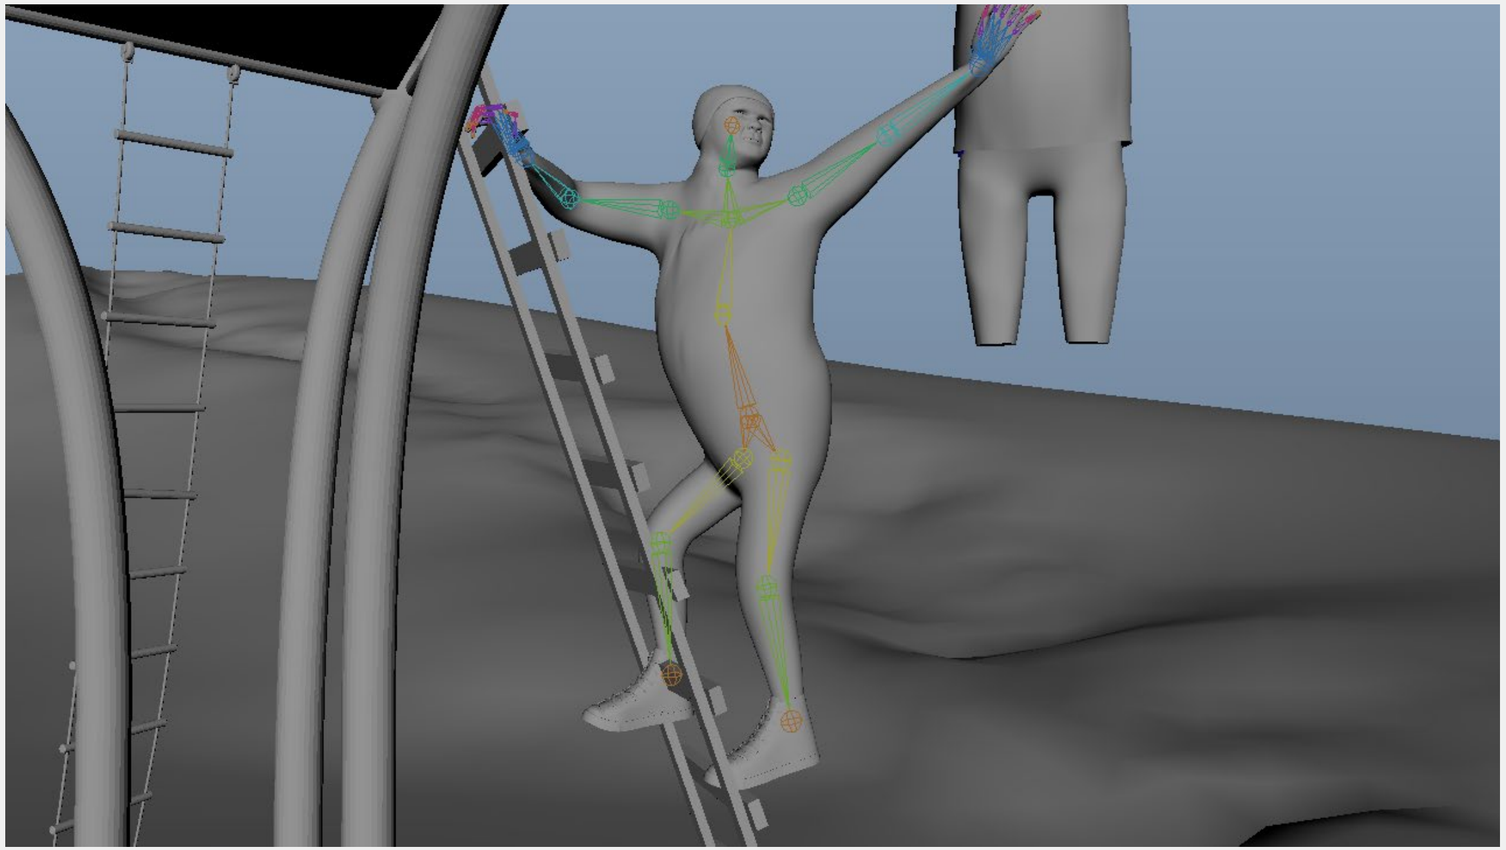

I created a simple rig for Dylan, with the extremities being his head, fingertips and ankles. This was all that was needed due to the model's simplicity.

I changed the weight paints on Dylan's legs a bit, so that moving one leg wouldn't affect the other.

Then I set keyframes for all joints at frame 1 forming a T-pose, then again at frame 100 for his climbing pose pictured here. In order for the nCloth to display correctly (which it isn't here), I needed to run a viewport simulation for every frame in-between so Maya could calculate the physics. The final render is from a few dozen frames later, so the cloth had time to settle after Dylan reached his destination.

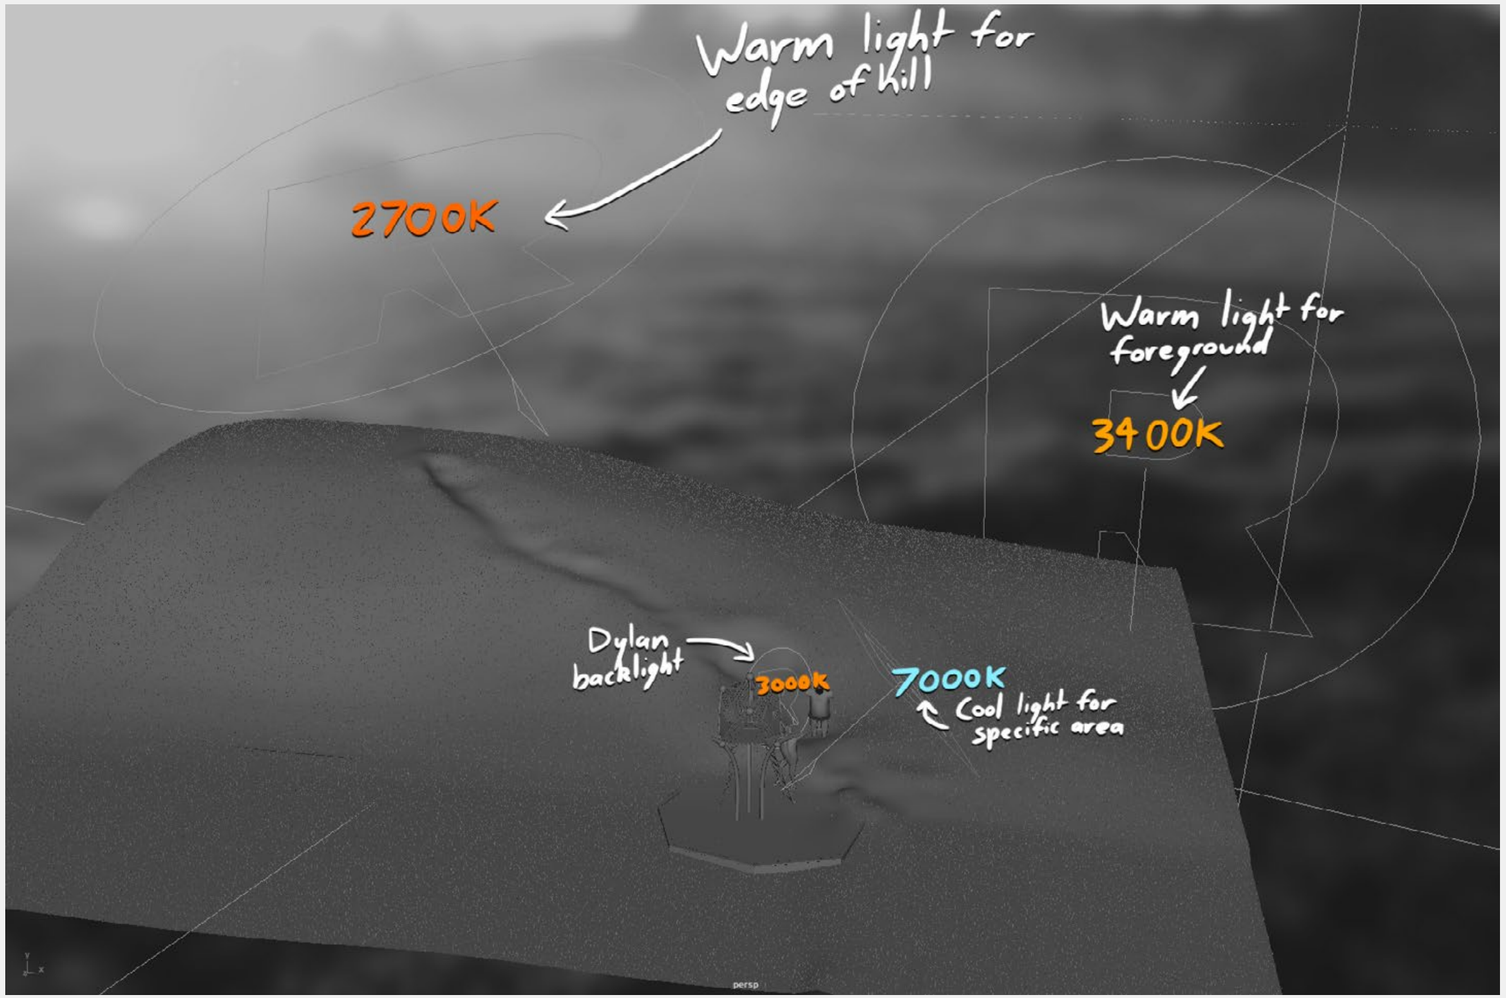

Lighting arrangements

Here you can see the lighting setup for the scene. There's 4 disk/ rectangle lights and a dome light with an HDR image applied (pictured at bottom). The dome light provides a background for the scene, similar to a skybox in a video game, with the added benefit of producing different colours to match the image. This results in nice realistic lighting.

I wanted different areas of the scene to look a specific way. For instance the far side of the hill would be lit with a golden glow, amplifying the light from the sun in the image.

A cooler light illuminates the treehouse in the foreground, to avoid going too far with the warmth. A warm backlight is positioned right behind Dylan, helping him feel more seamless with the environment.

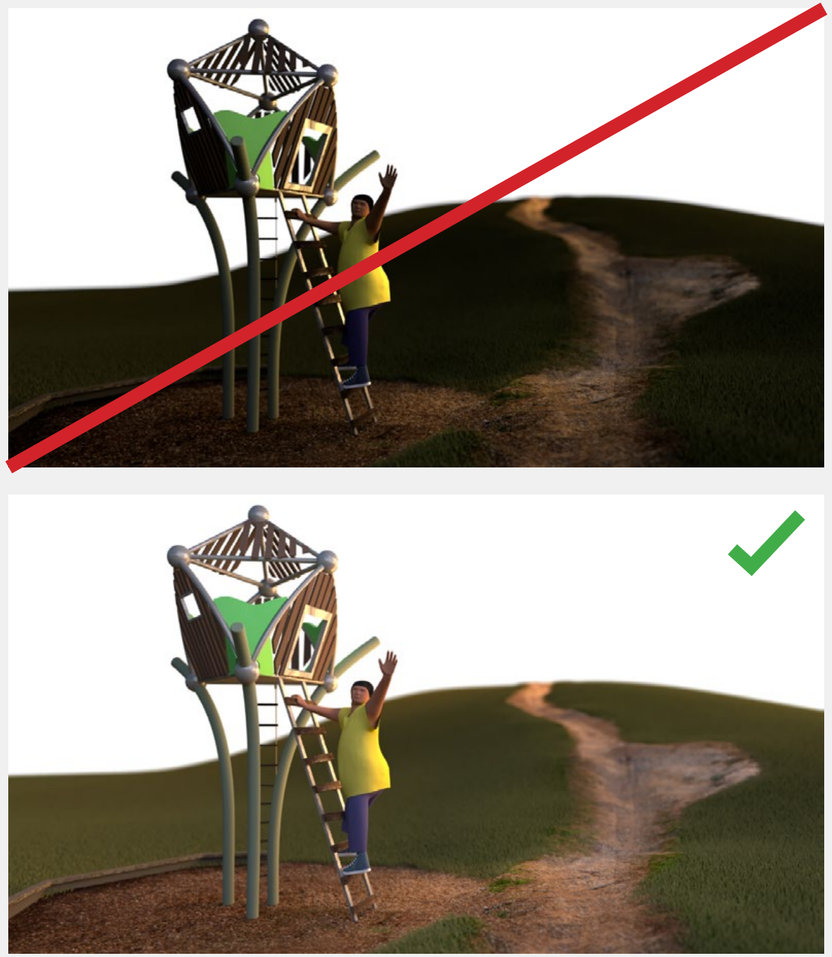

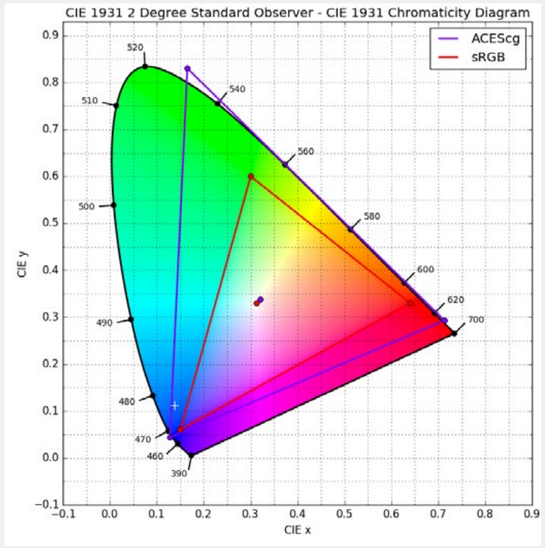

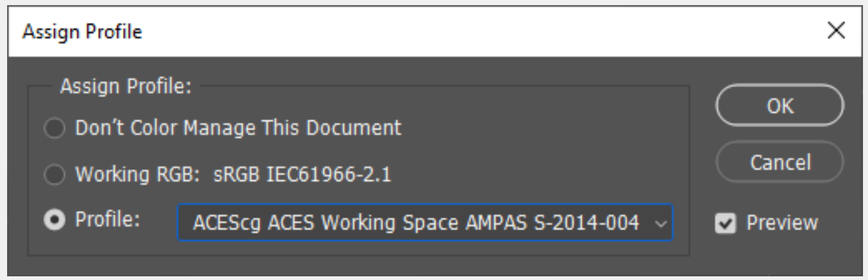

Assignment 2: CGI Character and Environment

Colour management

During this assignment I learnt about the ACEScg colour space. I don't claim to fully understand how it works, but the gist of it is it provides a wider range of colours than sRGB, the most common colour space.

Combining this with the scene being rendered in a 32bit EXR format meant the final exported render would contain vastly more light and colour information than an 8bit sRGB image.

Using this workflow properly meant bringing the rendered EXR into Photoshop, assigning the ACEScg profile from the Edit menu, then changing the image mode to 8bit, compressing everything down to the 8bit sRGB space.

My understanding is that ACEScg is more of a production format, when media is passed between various departments, in order to avoid limitations and compression.

Final render

And finally, here is the finished render.

You can see how the aforementioned light behind Dylan gives him a warm, inviting feel, combined with him waving to someone out of view.

Some textures are merely simple colours, but given enough time they would've looked similar in terms of quality to the ladder. If you look close you can see roughness on the metal beams, and specular areas on the wooden steps that respond to light. The grey balls surrounding the treehouse show the skybox in their reflections.

I never got around to creating trees and bushes, which is apparent as the scene looks rather empty and the background abruptly cuts to the skybox. I tried to mitigate this by introducing a shallow depth of field.