THE SURFACE OF ACCEPTABILITY IN VIRTUAL FACES

A Thesis

by

SCOT PHILIP ANDREASON

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Chair of Committee,	Frederic Parke
Committee Members,	Carol LaFayette
	Stephen Caffey
Head of Department,	Tim McLaughlin

August 2013

Major Subject: Visualization

ABSTRACT

This paper explores the surface properties of skin and eyes and their importance in the acceptance and success of a digital human face, specifically in relation to the uncanny valley. The uncanny valley hypothesis states that as a human representation approaches photo-realism, subtle differences from reality become unsettling. Recent studies suggest that the uncanny valley could exist over a far greater range, affecting abstract human representations as well. These competing findings are explored by analyzing how changes to the surface of a digital character affects it's level of acceptance. A female facial model is used as a base to compare a spectrum of different simulated real-world materials. The variations range from materials that are nearly identical to human skin, to those that are completely divergent from it, thus unnatural. After studying this catalogue of materials, it is concluded that given the right conditions, the uncanny valley can occur when facial representations are very near realism, as well as when human-likeness is quite distant from reality.

ACKNOWLEDGEMENTS

I would like to thank a number of people who helped guide and encourage me through the completion of this thesis. My chair, Dr. Frederic Parke, was helpful in developing the initial conception of this paper, and organizing my thoughts into a clear goal. He provided invaluable insight as a pioneer in the field of computer graphics. I appreciate my committee members, Carol LaFayette and Stephen Caffey, for providing alternative insight into the subject. I would also like to acknowledge the faculty and students at the Viz Lab at Texas A&M for discussions on the uncanny valley and developing my awareness in the field.

I want to thank Brandi Parish for her help with corrections, her artistic commentary, and her continued maintenance of my psychological state throughout the process. Thank you to my sister Sierra Andreason, and Jennifer Screen for being my photography subjects while gathering texture reference. And of course, I would like to thank my parents Mark and Lisa Andreason, whose support through the years allowed me to pursue this research, and whose not so subtle encouragement allowed me to finish it.

TABLE OF CONTENTS

Page

LIST OF FIGURES

1. INTRODUCTION

The drive to create increasingly realistic virtual characters in movies, games, and media has persisted since the advent of computer graphics (CG). Even now, as photo-realism is becoming more and more attainable, there exists a gap that prevents virtual characters from being completely believable as human. Overcoming this discrepancy by understanding where and why it appears is crucial to the creation of visually appealing digital humans.

In 1970, Japanese roboticist Masahiro Mori put forth the hypothesis of the "uncanny valley", which stated that as robots become more human-like they will eventually hit a threshold that makes them appear creepy and unacceptable (Mori, 1970). This idea has extended to computer graphics as the technology for rendering digital humans has advanced. More recent research in the field of human-computer interactions, conducted largely by Karl F. MacDorman, has found results indicative of a larger uncanny valley that can include human representations with less realistic characteristics (MacDorman, 2006; Green, MacDorman, Ho, & Vasudevan, 2008).

This thesis is a visual exploration to discover what surface properties exacerbate the uncanny valley, and what can be done to avoid the effect. The study evaluates the material properties of the skin and eyes on a digital face model and explores how the level of acceptance and appeal of the face changes as these properties are varied over a wide range of values. The skin and eye representations range from close approximations of accurate surface conditions, to abstract representations that deviate from realism considerably. This exploration allows the artist to define the points at which a face becomes unacceptable and even disturbing based on the specific surface qualities of the skin and eyes. This range is analyzed to discover whether the phenomenon only occurs very close to realism, or if it can also occur when the face is depicted with surface properties that are more unnatural.

Numerous studies have already probed the existence of the uncanny valley, exploring the theory through quantitative analyses based on viewers' reactions to different faces under various conditions. An admitted concern of past research was the lack of artistic judgment in the creation of visuals. Facial morphing and changes in the level of detail were all performed procedurally within a program (MacDorman et al., 2009). The virtual facial images encountered in movies and video games are not created using such an automated and aesthetically detached approach, but are instead artistically crafted, each displaying subtleties, nuances, and personal interpretations. By definition the uncanny valley appears close to realism. However, the digital faces used in most past studies were very unrealistic in their subtle geometric detail, shading and lighting quality, as well as rendering. Although many test images did appear unsettling, it is difficult to say whether this can be directly attributed to the uncanny valley. The current study explores the hypothesis qualitatively, from an artistic viewpoint. Each aspect of the final rendered images was painstakingly detailed to achieve the highest level of realism possible. The creation of the face progressed through a workflow similar to the practices common in a film or video game studio. As a result, many of the techniques used are directly relevant to industry practice. The successes and failures of the final images provide usable solutions and insight.

2. STATEMENT OF INTENT

The chief goal of this study is to explore how variations in the surface properties of a face contribute to perceived realism, and furthermore, what variations cause the greatest change in the appeal of the face. An additional goal of the study is to explore the face's range of personal acceptability to help determine where the uncanny valley occurs for the artist. Will the uncanny valley be evident when a face is close to realism, when a face is more abstract, or will the phenomenon occur at all?

Certain objectives must be completed to create a controlled study that will effectively explore the surface of the skin and eyes. The first objective is to create a facial model that is both aesthetically appealing and physically accurate. Special attention must be given to the subtle features that often detract from realism. Faces exhibiting more average traits (equated to beautiful) are perceived as comforting and therefore distract less from the focus of the study (Rhodes & Tremewan, 1996). Facial geometry that is unusual or unique will cause too much attention to be given to the form of the face, when it is the surface properties that are important.

The next objective is to create a lighting environment that allows a uniform view of the surfaces. The lighting should be dynamic and frame the form of the face in a way that is visually interesting, but not to the point that it becomes obtrusive (e.g. harsh and distracting). The background should be simple and not serve as a focal point. The third objective is to create a base surface for the skin and eyes that is as realistic as possible. The surfaces will simulate all natural light interactions such as subsurface scattering and accurate reflectance. The final objective is to create a library of skin and eye surface shaders that mimic other real-world materials. These surfaces range from materials that are similar to human skin, such as wax and silicone, to completely divergent materials, such as metal and glass. The eyes are studied in the same manner, but with different materials.

The skin and eye variations are studied as separate spectrums. While rendering the different versions of the skin, the eyes remain constant and use a base realistic eye shader. Likewise, the base realistic skin shader is kept constant while rendering the eyes. Combining variations of the skin and eyes adds a layer of complexity that this thesis will ignore. The base realistic skin and eye model are analyzed to determine their weaknesses and strengths at simulating reality, as well as their correlation with the uncanny valley. The material variations are explored to discover insights into the range of the uncanny valley. The findings will either aid the artist to support Mori's original hypothesis, to support MacDorman's notion of the wider uncanny valley, or find no evidence to support either.

3. BACKGROUND

The German psychologist Ernst Jentsch first explored the concept of "Uncanny" in his paper "On the Psychology of the Uncanny" in 1906. Uncanny arises from uncertainty towards the truth of what is being seen. Is a perceived living object actually alive, and is a perceived non-living object actually non-living? Discovering what causes this doubt could be key to avoiding the uncanny (Jentsch, 1906). Sigmund Freud expanded on the idea of uncanny. Freud speculated that the uncanny arose from the realization of one's own mortality from the failures of human replicas. Man copies himself to ensure immortality, but these copies only serve as a constant reminder of why they were created –we are going to eventually die. Also, the uncanny can be a result of confusion between fantasy and reality; when something imaginary becomes a fully realized, reality and is no longer symbolic. The uncanny is a concept characterized by fear, doubt, and death (Freud, 1955).

The hypothesis of the "uncanny valley" was first put forth by Japanese roboticist Masahiro Mori in 1970, in his article titled "Bukimi no Tani Gensho.'' Mori presented the idea that as robots become more human-like in their appearance and motion they will be increasingly appealing to humans, but when the robots reach a nearly human state there is a distinct drop in their appeal. Their subtle differences from humans will cause them to appear eerie and unsettling (Mori, 1970). Mori's theoretical graph (figure 1) illustrates the sharp decrease in appeal as human-likeness increases to a certain undefined point. Finding and defining this point has been the goal of numerous studies in the field of robotics. More recently, the uncanny valley has become a buzzword in the world of computer graphics as well, and has become a common criticism of virtual characters.

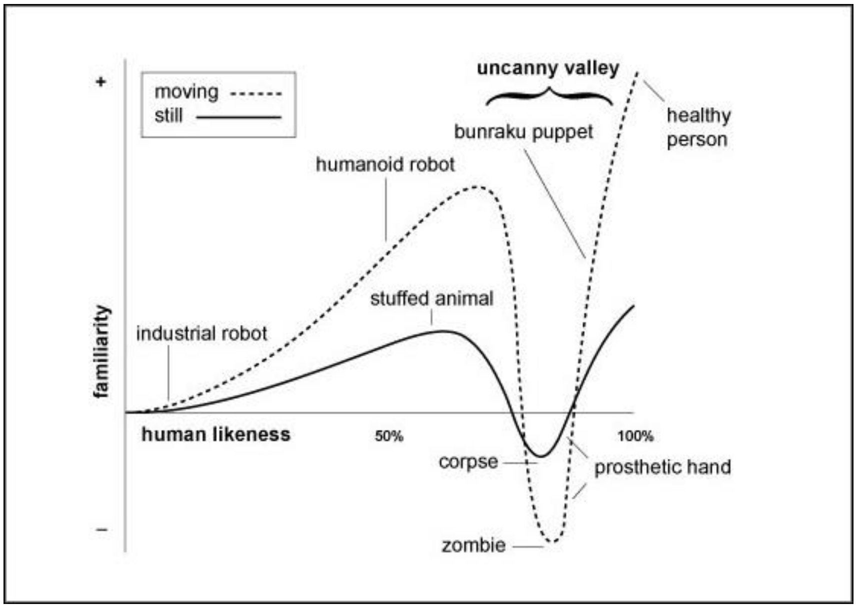

Figure 1. Graph of the uncanny valley proposed by Masahiro Mori.

Computer generated characters have existed as far back as the 80s, appearing in shorts such as "Tony de Peltrie", and Pixar's "Tin Toy". These characters exhibited rudimentary human form and motion that were clearly distant from reality. As technology advanced in the field, the level of realism in virtual characters increased, steadily approaching the uncanny valley. It wasn't until 2001, with the release of *Final Fantasy: The Spirits Within*, that the uncanny valley became a real topic of discussion in the CG realm (Kaba, 2013). This film marked the first attempt to create photo-realistically rendered characters utilizing full motion capture. The result was a relative consensus among critics that the technology was impressive, but the characters were distractingly creepy. Peter Travers of *Rolling Stone* commented "... you notice a coldness in the eyes, a mechanical quality in the movements" (Travers, 2001). It has been suggested that *Final Fantasy's* failure at the box office was due in large part to the uncanny valley, especially when compared to the success of the animated film *Shrek* that came out the same year. *Shrek* was highly stylized, with most characters intended to diverge from humans, thus limiting the uncanny valley effect. The same analysis can be applied to *The Polar Express*, which was completely reliant on motion capture and attempted photo-realism in the form and surface of its characters. The film suffered the same critical and box office shortcomings *Final Fantasy* had 3 years earlier. *The Incredibles*, which featured more abstracted character representations and exaggerated motion, was released the same year to critical acclaim and box office success (Pollick, 2009).

Even in very recent years, the uncanny valley has remained a clear barrier to the successful creation of photo-realistic characters. 2011's *Tron: Legacy* featured a digital recreation of a young Jeff Bridges, and became an exemplar of the uncanny valley. The same year saw the release of *The Adventures of Tintin: Secret of the Unicorn*, another feature length animated film that strived for hyperrealism. *The New York Times* proclaimed, "Neither fully human nor fully animated... this Tintin is lifelike, but without the pulse of real life" (Dargis, 2011). Although video games lag a bit behind films in terms of achieving a sense of visual realism, recent titles such as *Heavy Rain* and *L.A. Noire* have received heavy criticism for their depiction of uncanny faces. As the medium advances in rendering quality, many more cases of the uncanny will undoubtedly arise.

Numerous studies have probed both the existence of the uncanny valley, and the point at which it becomes apparent. Conditions such as eeriness are inherently hard to quantify, and as such, few studies have produced results that can pinpoint the drop-off in acceptance of digital characters. It is widely agreed that if the uncanny valley does exist, it is not quite as simple and elegant as the graph that Mori proposed. Competing evidence has shown that it is possible for the uncanny valley to occur over a much wider range of human-likeness. After presenting participants with a set of images that displayed the morphing of a robot head into a human head, MacDorman found the valley to exist at about 35-40% of human-likeness as opposed to Mori's theory of 80-85% (MacDorman, 2006). Again in 2008, a study measured responses to different human faces, robots, 2D and 3D human representations, and androids - all with varying proportions. The findings were inconsistent with Mori's original theory in this case as well. Subjects found faces the most creepy when they were rated as neutral on a scale of human-likeness, meaning these faces existed somewhere between human and non-human (Green, MacDorman, Ho, & Vasudevan, 2008). These findings suggest that the uncanny valley can exist over a far greater range than expected and that characters that are further from photo-realism can appear eerie as well. Similar experiments, where a head was morphed between a CG avatar and a photograph, revealed results that were more consistent with Mori, but only when the faces displayed abnormal traits such as unnatural eye shape. The valley is more apparent in characters that lack physical features commonly thought of as attractive and normal. Attractive characters are more likely to bridge the uncanny valley (Seyama & Nagayama, 2007).

Texture detail on the face can be even more important than form when conveying acceptable realism. Eeriness was much more widely reported when increasing the photorealism of textural details on the face than when increasing the geometric accuracy of the facial model. Furthermore, a mismatch in the level of texture detail on different elements of the face (e.g. photo-real skin and less detailed eyes) exacerbated this uncanny effect (MacDorman et al., 2009).

4. PROCESS

4.1. Methodology

An anatomically accurate face model with realistic detail was created as a base for the numerous material variations that were explored. The model was given a neutral facial posture that remained constant throughout the study. Realism was simulated in form, color and lighting, as well as reflectance and light interaction.

The skin and eye materials were first varied to represent surfaces with similar properties, including wax, silicone, rubber, and 'video game' (comparative to foam latex). The surface was then altered more noticeably to create very unnatural skin materials, such as porcelain, metal, glass, potato, and meat. These surreal variations were created to explore MacDorman's findings that suggest the existence of the uncanny valley over a broad range of human-likeness (MacDorman, 2006).

Variations in the eyes and skin were studied separately. The base eye shader was kept constant when studying the skin variations, and the base skin was kept constant when studying the eyes. A combination of static rendered images and turntables were generated to display the face with each material applied.

4.2. Face Model

The essential first step when creating the face model was to collect extensive reference for female facial anatomy. This reference came from Internet searches for high resolution images, personal photographs taken of volunteers, as well as classical drawings and paintings from the masters such as Leonardo da Vinci, Michelangelo Buonarotti, Auguste Rodin, etc. The face model was not based on any specific person, but was modeled as an amalgamation of facial features from numerous references. The traits of the face were constructed as an average of these references. People tend to prefer facial features that represent an average of the population, and typically consider more unique features unattractive. This predisposed view of beauty is referred to as the 'most average facial appearance effect' (Lidwell, Holden, & Butler, 2003, p.164-165).

The basic low-polygon face model was created using Autodesk Maya. The "box modeling" technique was used to begin blocking out the form. This method starts with a simple cube and builds in complexity as more edge loops and polygons are added. Vertices are pushed and pulled to shape the fundamental composition of the human face. At this point the face is an androgynous and hardly realistic rough structure. However, the polygons are arranged in a uniform and regular layout with flowing edge loops that allow the geometry to easily to be subdivided to create higher resolution detail later.

The base geometry was then imported into a 3D sculpting application called Zbrush. Working in Zbrush is more like working in virtual clay than moving vertices around and selecting polygons. This allows much easier creation and movement of large-scale forms and is essential for small details. In Zbrush, the shape and the relative placement of facial features were shifted considerably from the original base model. The head was then subdivided into layers of increasingly dense polygons, as more and more resolution was required to sculpt finer details. Once the form of the face was completely realized, small nuances such as wrinkles and pores were sculpted into the surface. At this point the face geometry had been subdivided 6 times and numbered in the millions of polygons. In that state, the mesh was far too dense to be re-imported and rendered in Maya. The model was imported back into Maya at its third subdivision level, which only consisted of a few thousand polygons. The remaining detail came from a displacement map exported from Zbrush that was used by the shader applied to the face. This geometry was subdivided at render time and the sub-pixels were shifted to represent the values in the displacement map (figure 2).

Figure 2. Face model geometry.

4.3. Eye Model

The eye model consists of two layers of geometry. The inner layer is a spherical polygon surface representing the sclera. The sphere has been indented in at its center to simulate the iris. The iris may not be completely accurate in form, but it will catch light in a way that appears convincing. Another slightly larger sphere was modeled around the sclera to represent the cornea. This sphere bulges out slightly at its pole, overlapping the indented iris. The tear duct in the corner of the eye was modeled as another separate piece of geometry. Lastly, a thin surface was modeled between the bottom eyelid and the eye. This geometry is transparent when rendered, and only catches specular reflections. It simulates wetness build up along the bottom eyelid edge.

4.4. Hair, Eyebrows, & Eyelashes

The initial plan was to use a bald head throughout the study to avoid the added complexity of hair. However, the absence of hair proved too distracting and unnatural after some preliminary test renders. The hair on top of the head was created using the "Maya Hair" system within Maya. This approach allows the placement of numerous "guide" hairs that will control the dynamic behavior and directionality of the actual hair curves seen only at render time. Each guide hair controls thousands of hair curves.

The placement of hair was determined based on 3D painting on the surface of the model using grayscale values. The value of 1 (white) indicates full hair, and 0 (black) indicates no hair. For the hair follicles to be distributed uniformly, they must be mapped onto a surface with consistently spaced UV coordinates. The proportional relationship between UV space and 3D space must be in near perfect alignment for the hair to be correctly distributed and avoid bald patches.

To achieve the look of hair pulled back in a ponytail, the guide hairs had to be dynamically blown backwards by using program forces such as wind. A constraint was placed on the back of the head to make the hairs converge and clump together at a point. This point simulates where the band would go to secure the ponytail. Since physical forces drive this process, there is limited control over hair placement. However, with a great deal of finesse, an adequate result was achieved. Once the hair was placed, numerous controls were used to adjust the look of the hair such as its spacing, curl, color randomization, and shading.

The eyebrows were created using "Maya Fur". This system is optimized for shorter hair that doesn't need to be dynamically placed. The hair can be combed into certain orientations. Again, there are numerous adjustments for randomization of color, spacing, direction, curl, etc. While there are slightly more controls available than for the hair system, an acceptable result required extensive work.

Creating the eyelashes was a painstaking process of modeling each follicle. NURBS cones were stretched and curved to correct proportions and placed along the eyelids. Eyelashes vary in thickness, length and direction from the center of the eyelid to the edges (figure 3). Exact precision and control was required.

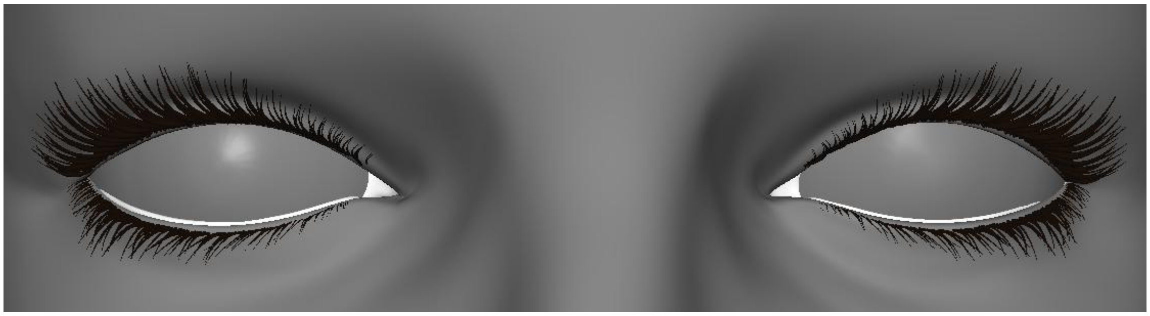

Figure 3. Eyelashes modeled and placed by hand.

Hair was rendered using Pixar's Renderman. Renderman is well integrated with Maya Hair and Maya Fur and is optimized for quick render times even when using self-shadowing and ambient occlusion on a per-hair basis. Since the face and eye geometry were rendered using a different renderer (V-Ray), the hair, eyebrows, and eyelashes were each rendered as separate passes composited together as a post process. The head geometry was rendered transparent while still receiving shadow information from the hair.

4.5. Lighting and Rendering Setup

The lighting configuration was purposely kept simple and non-obtrusive. It is a basic three-point lighting setup with a relatively strong key light. The fill light is approximately one third the intensity of the key light, enough to keep the shadow side of the face from appearing too dark while keeping good light to dark contrast. The rim light is about three times brighter than the key and is intended to add interest to the silhouette of the head, as well as intensify the sub-surface effect on the thinner regions of the face, predominantly the ears.

Overall the lighting is neutral and uniform, with enough shadow information so the head doesn't appear flat. The lighting is intended to appear realistic without being overly dramatic. The lighting was designed to focus attention on the surface properties of the face; harsh shadows or strong highlights were avoided for this reason. Large area lights were used to cast soft shadows. Global illumination was employed for secondary lighting.

To create global illumination, the face model was contained within a spherical environment that allowed light to bounce around and illuminate indirectly lit areas of the model (figure 4). Global illumination was achieved using an irradiance map for the primary indirect lighting calculation, and a photon map as a secondary calculation. The spherical environment is mapped with an HDR image to create realistic reflections on the geometry. The skin and eyes were rendered using the V-Ray physically based renderer.

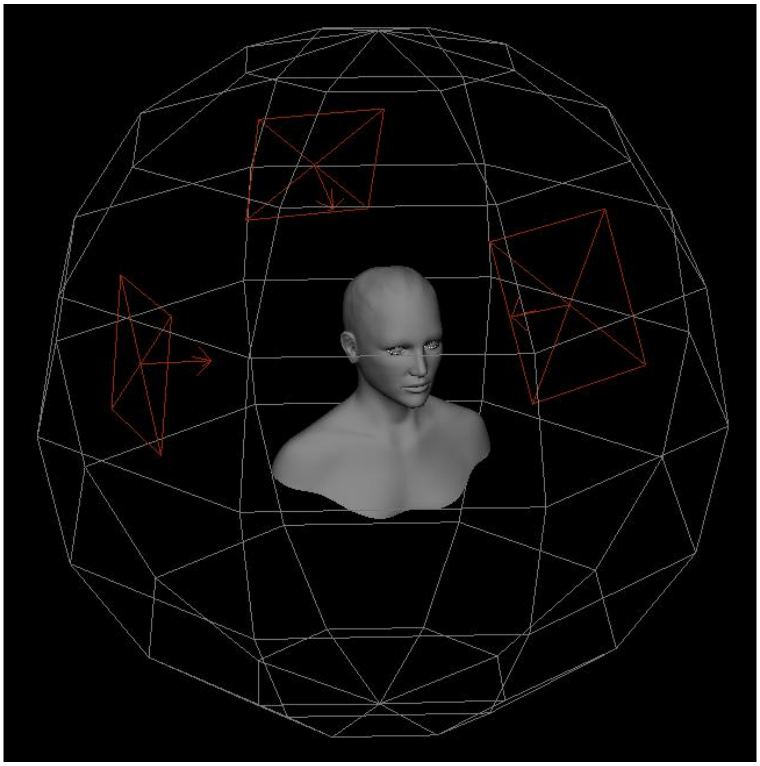

Figure 4. Lighting setup.

4.6. Base Skin Texture

Diffuse color is perceived as the pure color of a surface under white light without any reflection. For this project a diffuse color map was created from high definition photographs taken of human faces. Three separate photography sessions were held with three different volunteers who were photographed under studio lighting conditions. The live models were surrounded by three large lights equipped with diffusers. Direct light was avoided so bright reflections from oil on the skin would not be present in the images. Reflection highlights were absent from the color map since they were added later as a specular layer in the surface shader. Photographs were taken from a front-on view, a direct side view, a three-quarter view, an under chin view, and numerous close-up views of high detail regions such as the ears, lips and nose.

The face geometry was imported into Autodesk Mudbox to perform basic image projections from the photographic reference taken. Mudbox is a 3D sculpting and painting program with features and functions comparable to Zbrush. Mudbox was used in order to take advantage of its improved texturing capabilities. The texture resolution of images projected in Mudbox is limited only by the resolution of the original photographs.

The process of projecting an image is one of the most basic forms of texture painting: it involves placing a 2D image over 3D geometry and transferring the pixels of the image directly to the UV map of the 3D geometry. However, the projection will only be successful from the exact perspective of the image and will appear stretched when the geometry is viewed from any other angle. For this reason, numerous projections must be made to account for all viewing angles. This can be a very time-intensive process that requires images taken from all angles of the face. In this case, only a few projections were made, one from the front, one from the side, and one three quarter. These rough projections were made to block in the main features of the face and to easily match their location with the UVs of the face geometry.

The main texture work was done in Photoshop. The projections from Mudbox didn't match up with one another along the edges and were stretched in areas of high curvature such as the nostrils, ears, and around the nose. The fix work in Photoshop was a time-intensive process of filling in these areas with patches of texture information from the photographic images and stitching them together seamlessly with the projections. The result was a unified image without any areas of stretched pixels and a uniform high-resolution level of detail. Color was overlaid on some areas of the face to enhance subtle variations in the skin that had been lost in the images. Finally, a faint layer of makeup was painted to enhance the beauty of the skin and for added contrast. The result is shown in Figure 5.

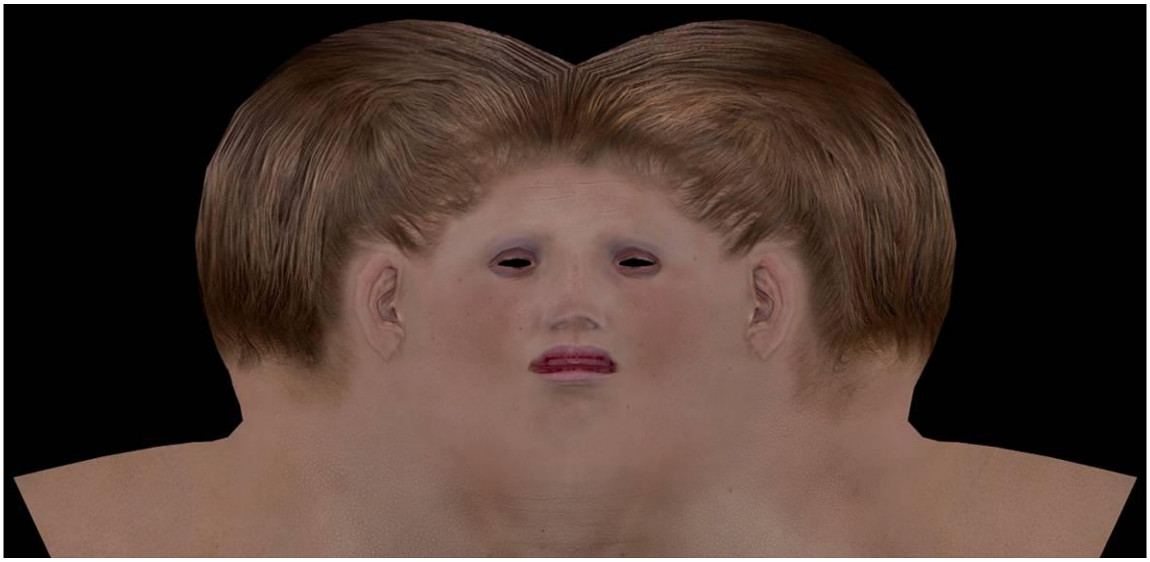

Figure 5. Skin texture map layout.

A sub-dermal map was created based on the diffuse color map. The sub-dermal map is used to simulate the tissue under the skin. It is used in determining the color of sub-surface scattering at each pixel. This map was created by first level-adjusting and saturating the diffuse color map, and then tinting the values red. Other details such as veins and color variations of blues and purples were painted on the map by hand.

The specular reflection map, which determines the intensity of light reflection from the surface, was also a derivative of the base color map. The color map was de-saturated and level adjusted to bring out small details in the pores and wrinkles. Oily areas of the face such as the cheeks, nose, and forehead were brightened to catch more reflections. The lips were brightened to near white to ensure a stronger reflection.

A roughness map, which looks very similar to the specular map, was also made to define glossy or matte values for the specular reflections. The areas that were made slightly more reflective in the specular map were increased in gloss value in the roughness map.

Changes in roughness greatly influence the regional reflection properties of the skin. It can also be used to make the lips appear wet. The roughness map operates on grayscale values from 0-1, with 1 being mirror-like glossiness.

4.7. Skin Shaders

4.7.1. Base Skin Shader

Figure 6. Base skin shader.

The base skin shader (figure 6) is composed of a reflection layer, a diffuse layer, and a subsurface scattering layer. The diffuse color layer has only a small contribution, with most of the final color coming from the subsurface scattering layer. The light scattering depth within the skin is set relatively low (0.2cm). This number was not based on any real-world calculations, but was arbitrary, based on the unit size of the mesh and the intensity of the light to produce a visually acceptable approximation of reality.

Most water-based materials such as skin exhibit strong forward scattering, meaning that light will scatter mostly forward (in the same direction it comes from). Before light leaves the skin, it is reflected many times inside. This phenomenon is referred to as multiple scattering. In this case, light will bounce around many times before it leaves the surface, as opposed to single scattering where light is reflected only once. Skin exhibits both single and multiple scattering, but multiple scattering is much more predominant. Multiple scattering changes the direction of the incoming light rays numerous times before it exits at an undetermined point (Igarashi, Nishino, & Nayar, 2008). The result is that the direction of incoming light does not influence the subsurface effect, creating a random diffusion of light.

In the V-ray renderer, scattering direction is controlled by the “phase function,” which ranges from -1.0 to 1.0. In this case, the phase function was set to 0.8 to simulate strong forward scattering and increase the effect of multiple scattering. The sub-dermal map controls the sub-surface coloration.

Reflection from the skin is not very strong or sharp, resulting in muted reflections. The skin reflection layer uses a Blinn shading model, with the amount of reflection being driven by a texture map. The Blinn shading model is an algorithm that defines how light interacts with a surface, and controls how diffuse or reflective the surface will appear. The skin has a slight fresnel effect that is dependent on the index of refraction. The fresnel effect is reflected light that can only be seen at the glancing angle of a surface and is dependent on the point of view of the camera (Birn, 2000). The index of refraction (IOR) of light in the skin is about 1.45 (Wilhelm et. al., 1997). The light reflections seen on the skin are the result of a spherical HDR map that wraps around the environment, but is hidden in the final render

4.7.2. Wax Skin Shader

Figure 7. Wax skin shader.

The wax shader (figure 7) relies on a deeper scatter radius than the base skin shader to achieve the overall effect. The rays penetrate much further into the surface and bounce around before exiting. This gives the effect of a less diffuse and more translucent surface. The scattering depth is set to 1cm, or five times deeper than the base skin model. Similar to the base shader, the wax exhibits forward scattering with predominately multiple scattering effects under the surface, with the 'phase function' set to 0.8. There is no blood under the surface in this case, so the sub-dermal color is set to a more muted pale orange.

The wax has slightly stronger reflective properties than the base, with a noticeably sharper highlight. The reflection map, which controls changes in reflection properties across different areas of the face, has very low contrast in this case. Together, this results in a surface that has a uniform, slick, smooth appearance. The fresnel effect and IOR are unchanged from the base.

4.7.3. *Video Game Shader*

Figure 8. Video game skin shader.

The video game shader (figure 8) is rather unique since it is the only shader that does not simulate a real-world material. Instead, this shader attempts to replicate the material properties that would appear on a character from a current generation console, such as "Playstation 3" or "Xbox 360". A human in a video game must be rendered in real-time and must forgo some of the more intensive shading algorithms that offline rendering can afford.

The most recognizable difference in this shader (as compared to other shaders) is the absence of sub-surface scattering, which results in the skin appearing as a flat, solid surface.

The diffuse shading becomes most evident in areas with high curvature, specifically the ears and lips. Sub-surface scattering would help to lighten the overly dark shadows in these areas.

The changes in specular intensity and roughness are not very noticeable, showing fairly uniform light reflectance over the surface. There is also no fresnel reflection present, so no changes in specularity at glancing angles. The resolution of all maps used in the shader have been reduced to half-size, to simulate a gaming environment that utilizes lower resolution maps to increase rendering efficiency.

4.7.4. Rubber Shader

Figure 9. Rubber skin shader.

The rubber shader (figure 9) has a relatively deep scatter depth of 1cm. However, its phase function is set at -0.5, meaning the light will mainly scatter backwards in the direction it came from. This combination of deep scatter depth and backwards scattering gives the effect of a strong solid surface, with some malleability. The sub-dermal map has a blue-purple tint to give the skin a more synthetic, non-living look.

The IOR of rubber is set to its real-world vale of 1.519 (Wood & Tilton, 1949). The level of reflectance is relatively high to simulate the qualities of the rubber found on a Halloween mask. The specular map is also tinted orange to create a hue shift when combined with the bluish sub-dermal map. This effectively kills the variation of color found in the base face shader.

4.7.5. Silicone Shader

Figure 10. Silicone skin shader.

The silicone shader (figure 10) is similar to the base shader. It is essentially a pristine version of the base, devoid of any imperfections. The scattering depth is double the base at

0.4cm. This creates a softer appearance as the light travels deeper into the skin and diffuses. The index of refraction is set at 1.2, noticeably lower than the base skin (Li, 1980). The phase function is set at 0.8, relying on forward multiple scattering like the base shader. The color map has been airbrushed to get rid of fine details such as pores, imperfections and skin discoloration. It has also been colorized to appear a uniform pinkish-white. The color of the lips and the makeup have been exaggerated to stand out. The displacement map and specular map have also been touched up from the base to hide any high-frequency detail that would be seen as imperfections. The altered specular map makes the surface appear much smoother, and a bit more glossy. However, the actual specular properties are identical to the base skin.

4.7.6. Porcelain Shader

Figure 11. Porcelain skin shader.

The porcelain shader (figure 11) represents a hard, solid material with very little light passing through the surface. For this reason the scattering depth is very shallow at 0.05cm. The index of refraction is 1.6, and the phase function is -0.7, meaning almost all light rays are reflected back along the direction they came. This is accurate for a dense material such as porcelain (Ralph & Chau, 1993). A smooth airbrushed color map, like the one used for the silicone shader, is used in this case as well. The flesh tones have been muted substantially, and the entire map has a white overlay. The final porcelain color map is actually not much lighter than the map used for silicone, but the absence of subsurface scattering gives the appearance of a much more stark white surface.

The specular map is also smoothed out so there is little variation in reflection properties over the surface of the model. The surface is very glossy, with well-defined reflections of the surrounding environment. The fresnel highlights are very noticeable at the edges of the surface, insuring that the material stays reflective and doesn't darken at glancing angles. The hard specular highlights help to define a rigid, inorganic material that would feel slick to the touch, and completely solid in form.

4.7.7. Potato Shader

Figure 12. Potato skin shader.

The potato shader (figure 12) is the first material that uses a completely different color map from the base shader. Other shaders such as porcelain and silicone have manipulated the color map, but this is the first shader that is a complete deviation. This material marks a transition into more abstract surface representations to follow. The color map is taken from photographs of a potato at different angles that are then stitched together in Photoshop. The displacement map and specular map are derivatives of the color map that use the same information from the photographs.

A potato has a shallow translucency, but the overall subsurface scattering effect is minimal. In this case the scatter depth is 0.2cm and the phase function is 0.25, demonstrating a shallow light penetration that scatters a bit more forwards than backwards. Overall the surface will appear opaque except in very thin areas such as the ear of the model.

The specular map has been drastically altered from the base to match the contours and imperfections of a potato. The surface of the potato is rather rough, and highly diffuse. The specular contribution is low, producing very diffuse reflections, and no Fresnel effect.

4.7.8. *Meat Shader*

Figure 13. Meat skin shader.

The meat shader (figure 13) is another very unnatural surface variation. This material is not meant to represent human muscle tissue under the skin, but rather a slab of raw beef steak. Like the potato shader, this unnatural deviation uses a unique color map that is different from the base skin texture. In this case, the color map is taken from photographic reference of steak from a butcher's block. The displacement map is derived from the color map and slightly blurred to reduced noise. The displacement map is used to show differences in elevation of the bits of fat, connective tissue, and the porous surface of the muscle.

The subsurface scattering values used in the meat shader are taken from personal approximations after studying the light interactions with a slab of raw steak. The meat appears similar to skin, but more dense. The scatter radius is 0.2cm (same as base skin), but the phase function is 0.5, making the surface reflect more light rays than the skin shader. With slightly less rays scattering inside the surface, the sub-surface effect is less noticeable than on the base skin. Also, the sub-dermal map used takes into account the areas of fat, which are less dense than muscle, and allow more light penetration.

The meat is meant to appear moist with some blood still under the surface. Therefore, the specular reflections read as very glossy. The intense localized reflections make the meat shader appear to approximate a wet surface. The specular map is based off of the color map to show the changes in contour and tissue type in the meat. Most of the specular breakup is made apparent by the detail in the displacement map. There is a slight fresnel contribution to make the meat appear wet on the edges.

4.7.9. *Metal Shader*

Figure 14. Metal skin shader.

The metal shader (figure 14) represents a tarnished aluminum surface. The diffuse color map is the same as the base shader, but its contribution is very low, which makes the color appear muted. Metal materials rely mainly on specular light interaction for the final color of the surface; otherwise they will appear more like plastic than metal. Also, metal is completely opaque so there is no sub-surface scattering. The shader is almost completely dependent on the specular contribution. The specular map has some random variation in value to add interest to the surface. The specular map helps to break up the smooth reflections that would otherwise cause the surface to appear more like chrome. The reflections are intense and well defined, and there is a noticeable fresnel contribution around the silhouette of the surface.

4.7.10. Glass Shader

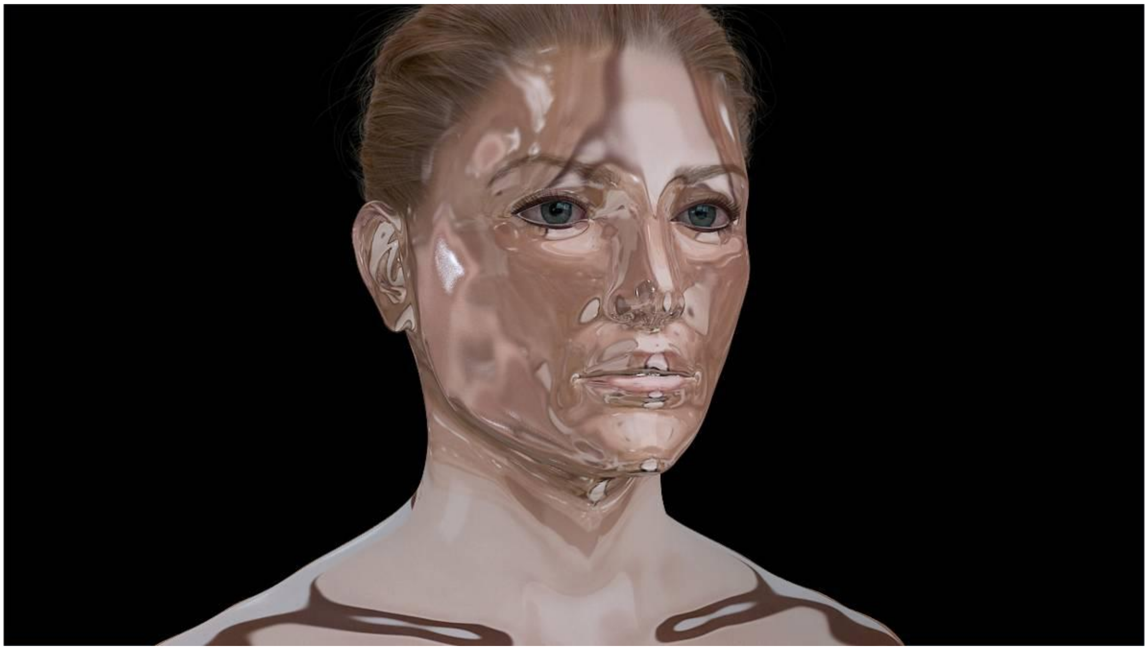

Figure 15. Glass skin shader.

The glass shader (figure 15) simulates thick semi-transparent glass. The base color map is used in this case, but the transparent nature of the glass causes the colors to appear very faint and only tint the surface. The look of glass is mainly achieved through light refraction. Standard glass has an index of refraction around 1.5 (Malitson, 1965). In this case the IOR has been exaggerated to 1.9, to accentuate the bending of light for a more interesting image. There is a "fog" parameter within the refraction layer that allows for a more foggy semi-transparent glass representation. This addition makes the glass appear thicker and helps to show the volume of the mesh better. The face is reflecting and refracting its own geometry, as well as an HRD environment map. The reflections are very sharp with no breakup, to show a smooth, perfect surface.

4.8. Eye Shaders

4.8.1. Base Eye Shader

Figure 16. Base eye shader.

The base eye (figure 16) consists of two different shaders, one for the sclera geometry, and one for the cornea. The sclera is mainly dependent on a diffuse color map. There is a slight sub-surface contribution that uses the same values as the base skin. The specular map limits reflections to only the iris to get a faint glint of light. The specular map on the iris is deeply saturated green so it has a less metallic white highlight.

The cornea shader is the main specular component of the eye. It is completely transparent, so there is no diffuse contribution. The reflections on the surface are almost perfectly defined with very little blur. The reflection glossiness is set at 0.99, simulating mirror-like reflection clarity. This gives the appearance of a wet surface. There is a very slight noise bump map used on the cornea to subtly show some undulation on the surface and slightly shift the reflections. The shader reflects an environment image mapped on a sphere surrounding the scene, as well as the face geometry. The eyelids and eyelashes can be seen as faint reflections. The cornea has an index of refraction of 1.4, which causes the sclera underneath to be refracted and appear slightly larger (Patel, Marshall, & Fitzke, 1995).

4.8.2. Plastic Eye Shader

Figure 17. Plastic eye shader.

The plastic eye (figure 17) is a simplified version of the base eye. In the color map the veins and redness around the edges are removed and color variation in the iris has been greatly muted and blurred. The sclera has no subsurface scattering, but otherwise its shading is the same as the base. The cornea shader has less defined and blurrier reflections than the base.

Also there is a fresnel contribution to give a plastic look. The index of refraction has been set unrealistically high at 3.5 to make the iris and pupil appear larger. This gives the appearance of more cartoony stylized eyes.

*4.8.*3. *Foggy Glass Eye Shader*

Figure 18. Foggy glass eye shader.

The foggy glass eye (figure 18) uses the same material properties as the base eye for the sclera layer, with the noted absence of sub-surface scattering. The most noticeable differences come from the cornea shader. Much like the glass skin shader, the refraction layer includes a fog parameter that makes the cornea appear slightly thicker and partially occludes the sclera. This gives the appearance of glass with a dirty film over the surface. The reflections are more diffuse than the base eye, making the surface appear drier and slightly impure. The IOR is 1.6, which is within the range of physical glass (Malitson, 1965).

4.8.4. Rubber Eye Shader

Figure 19. Rubber eye shader.

The rubber eye (figure 19) only uses the cornea geometry. The outer cornea has a completely opaque diffuse color shader applied to it, with no contributing shader for reflections. The rubber material is actually very simple. Most of its information comes from the diffuse layer and is based on a simplified color map similar to the plastic eye shader. As is the case with the plastic eye shader, the pupils are slightly larger relative to the iris to enhance its unnatural appearance. The specular contribution is faint with very diffuse reflections. The reflections are subtle to contrast with the other eye shaders, especially the realistic human eye.

*4.8.*5. *Gelatin Eye Shader*

Figure 20. Gelatin eye shader.

The gelatin eye (figure 20) is the most complex and experimental of the shaders. The sclera material is very dependent on sub-surface scattering, with a scattering depth of 0.5 cm and a phase function of 0.8. The shader mimics very translucent gelatin, allowing light to travel fairly deep into the surface. The result is a surface that appears less solid than the base eye. The base color map is the same simplified base map used for the plastic eye shader, highly saturated to mimic the bright food dyes used in colored gelatin.

The cornea shader has a bump map with subtle noise to help break up the specular reflection and simulate imperfections in the gelatin surface. The reflection contribution is 1.0, effectively reflecting all light from the surrounding environment, and making the specular highlight very bright. The reflections are very glossy to suggest a wet, slippery surface. There is a relatively strong fresnel contribution to keep the gelatin material reflective along the edges.

*4.8.*6. *Metal Eye Shader*

Figure 21. Metal eye shader.

The metal eye (figure 21) only uses the cornea geometry, similar to the rubber eye shader. There aren't two separate shaders for the sclera and cornea, just one for the cornea that is completely opaque. The base color map is used, unchanged from the original. However, the diffuse contribution is very low since metal color is mainly a result of specular reflection. The metal is highly polished with some slight variation in roughness coming from a grayscale map. The shader uses anisotropic reflections for a more directional specular highlight, simulating a slightly brushed metal surface. The IOR is 2.5, comparable to a highly reflective steel surface (Sugimoto & Matsuda, 1983).

5. INFORMATION AQUIRED

5.1. Better Skin by Comparison

At certain points during the creation of realistic skin, the face appeared very similar to some of the later variations, such as wax and the video game shader. To arrive at a successful solution, it was the natural course of action to experiment with parameters pushed to high and low extremes. Even once the realistic skin shader was initially completed, there were numerous times it had to be revisited throughout the process of making other material variations. For example, the base skin shader appeared far too dry and lacking in sub-surface scattering when initially compared to the wax shader. Without this comparison available, the base skin would have appeared more like the video game shader. It became easier to simulate different materials when viewing them in direct comparison to a similar material. It was very obvious what needed to be altered to make a material unique and different, when considered side by side with another comparable material.

The wax and video game shaders can be viewed as two examples of mistakes made when attempting to create realistic skin. Rubber and silicone (when ignoring color and analyzing shading properties only) are good examples of the same problems, but less extreme. One may find it easy to use sub-surface scattering as the instant solution to give the skin depth and the appearance of life, but the result can be similar to the wax shader when the effect is pushed too far. Even in the most technologically advanced feature films of today, sub-surface scattering can be noticeable: characters' skin appears too soft and translucent, and more malleable than firm.

A problem present in the wax shader is the oily appearance of the surface of the skin. The skin naturally produces oil, but its appearance is less uniform throughout. Areas such as the

"T zone" of the face usually appear more oily, and imperfections and pores on the face break up the smooth, uniform glossiness. The video game shader pushes the representation of skin too far in the exact opposite direction. Here sub-surface scattering is absent and there is almost no oil on the surface of the skin. Overall the skin appears hard and dry with little variation in color value.

One may find that shaders can make the color of the skin appear too uniform and lose variation. Exaggerating what appear to be subtle color shifts is key to achieving a more realistic final render. For this thesis, this effect was accomplished mainly by pushing hue and saturation to the extreme in the sub-dermal map. This is an important change that was made to the base skin shader to set it apart from the color uniformity of the video game shader. The common inclination to avoid too much variation in natural skin color shifts, pores, and imperfections when creating a young female face results in an abnormally pristine surface akin to the silicone shader.

5.2. Creating Realistic Eyes

The eyes are the feature of the face perhaps most crucial to communicating life (O'Neill, 2008, p. 13). They are also a difficult feature to reproduce realistically. People spend a great deal of time studying each other's eyes, and have subconsciously become experts at spotting when something is incorrect about them. This natural familiarity poses a great challenge when creating CG eyes. This is also true for the rest of the face as well, and to a slightly lesser extent, the entire human body. However, considering the complexity of the eyes, composed of numerous layers and a liquid film on the surface, the difficulty in achieving a successful reproduction can be even more appreciated (Stahlberg, 2007, p. 92). Also, research has found that a human recognizes eyes slightly slower than they recognize the face (Watanabe, 1999). The final burden of believability falls on the eyes: if the face is deemed acceptable, but the eyes discredit it, the uncanny valley factor can be extreme.

This study demonstrates how slight changes in shader properties and coloration can cause the eye to be instantly recognizable as fake. One of the most important, and the most difficult feature of the eye to achieve correctly, is accurate reflectivity. Numerous considerations must be made about reflections, including glossiness, subtle changes in the surface of the cornea that affect the shape, tear build-up, the placement of the reflections, and the reflected image itself.

Giving the eye the appearance of wetness is key to achieving realistic eye representation. The differences between the plastic eye and the base realistic eye are actually rather slight. Aside from the more simplified color map, a small change in reflection glossiness is the only other main variance. The reflection is slightly more diffuse, causing the eye to appear plastic. Pushing this effect much further results in a look similar to the rubber eye, which may arguably be the most disturbing of the eye variations. Accurate reflections (as well as shadows) also help create a smooth transition between the eyelid and the eye. The eye reflects the eyelids and eye lashes to create darkening around the edges of the eyes. This, along with the tear layer at the base of the eye, helps soften the harsh contrast from eye white to skin. This effect is very subtle, but it is noticed when missing, as in the case of the rubber eye. There is no transition around the edges, causing the eye to appear as a separate element stuck on like a sticker.

One important key to a believable reflection highlight is a good HDR environment image with a defined region of bright light. A uniformly bright map, or a map with multiple strong light sources, will result in a muddled reflection that makes the surface appear plastic. The HDR environment image should have a limited number of strong light sources. For example, in this study the HDR image uses a window light source, utilizing the window frame to define the area of light. It is helpful to breakup and define the high light areas to add interest and realism. Also, it is just as important to make sure the most intense part of the reflection appears on the correct region of the eye. A reflection falling on the iris, but not completely occluding the pupil, creates a strong contrast and is the best location. An intense reflection to one of the far corners of the eye appears distracting and does not pop out against the white of the eye (figure 22).

Figure 22. Eye Reflections.

The physical surface of the cornea also affects the reflection. In reality no organic surface is perfectly smooth, and this is true for the cornea as well. The corneal epithelium, the outermost layer of the cornea, has irregularities that disrupt the film of tears on its surface (Dohlman, 1971). This can easily be achieved in CG by adding very subtle noise into the displacement or bump component of the shader. This will slightly shift the reflection and suggest very small imperfections. This effect should be present, but not noticeable.

6. RESULTS

The results of this study are subjective, and can be attributed to my own personal opinions due to various perceptions and emotional responses created by the images. The uncanny valley is not a scientifically proven fact, and there is variance from person to person concerning what about an image is unsettling and at what point it becomes unsettling. The exploration and analysis are based on personal conclusions and conjecture. The findings are not quantifiable, but can be viewed as an artistic exposition of the competing theories of the uncanny valley seen across a range of images. The different skin and eye variations represent two theoretical sides of the uncanny valley; one close to realism, and the other distant and unnatural. The following analysis will center on where the biggest shifts in personal acceptance of the face occur and what factors contribute to the uncanny response.

The base "realistic" face model is obviously not a perfect representation of the human face. It would not be difficult for the average viewer to spot that this face is not a photograph but a CG model. There are many aspects that cause the face to fail, some that cannot even be fully articulated but that are clearly present at a subconscious level. Our natural ability to spot an imposter, or something inhuman attempting to appear human, tells the viewer that this face is not real. However, the real question is: does this cause the face to appear uncanny?

The uncanny valley hypothesis does not simply state that any face unsuccessful in achieving complete realism is uncanny, but rather there is an undefined threshold near photo-realism that causes a visual (and thus psychological) disturbance. Even though this face model doesn't achieve reality, it does not necessarily fall into the uncanny valley. It could even be argued that this face model is able to avoid the uncanny valley completely.

6.1. Why the Face is Successful

Although the face is not a perfect photo-realistic human representation, it would be fair to state that the face comes fairly close to achieving realism. One could easily perceive that the intent of the image is a truthful representation of a human female face, and that the final result is a largely accurate depiction. It could be argued that this image, which is very close to realism, is able to avoid falling into the uncanny valley. A question remains as to how, exactly, a face that fits the description of Mori's uncanny valley able to bridge its gap.

I believe the success of the base model is largely due to the consistency of detail in its elements. The uncanny effect becomes noticeable when a single part, or parts, of a face stand out as discrepant from the combined face as a whole. It could be argued that a CG face is only as successful as the difference between its strongest and weakest aspect. A noticeable separation in the believability level of different features of the face seem to directly contribute to how unsettling it appears. Perfectly rendered skin can be completely undermined by a glassy looking eye. An impeccably rendered hair model next to an inferior skin shader will make the skin look that much worse. In these cases, the weaker aspects of the face could have been much more successful had the comparison not been made between it and the higher quality aspect. Absent this disconnect, all parts will combine to make a cohesive whole. This is not to say, however, that if the skin is easy to render realistically but the eye is troublesome, one should downgrade the skin quality to match the eye. Keen artistic decisions are the key to ensuring that all elements meet at a mutually acceptable level of realism that minimizes any separation in believability.

The base model displays skin, eyes, hair, and geometry that are all at a similar level of detail. No part instantly stands out as exceptionally well rendered or unusually poorly rendered.

This is not saying the model qualifies as mediocre; but rather that it works as a unified whole by not causing any single area to stand out as higher or lower quality. In reality, a person doesn't look at another person's face and think how successful it is at being a real face. However, a CG face is subject to constant critical scrutiny by its audience. Provoking the viewer’s critical assessment is not the intended goal of the CG face, which is created to help tell a story or facilitate some interaction. Not providing the viewer with any areas to pick out and critically dissect can help obviate the viewer’s need to compare the face and its component parts to a real human face encountered in the real world. In addition to the rendering quality, the base face includes the aesthetically appealing advantages of a calm neutral facial expression and a reasonably attractive form.

6.2. Where is the Uncanny Valley Present?

The main goal of this study was to explore where the uncanny valley occurs, and what factors lead to the phenomenon by studying a spectrum of skin and eye representations. Furthermore, this exploration attempted to determine if the uncanny valley is more apparent in cases when the skin and eyes are close to reality, or when they are more distant and surreal.

Figure 23. Skin and eye material catalogue.

6.2.1. The Skin

The various images and videos created during the implementation of this thesis (figure 23) were studied and broken down to determine what aspects contribute to the uncanny valley. The occurrence of the uncanny is clearly the sum of numerous factors and is not easily defined at any specific point. The faces were viewed side by side to more easily examine them based on their relationship to one another. They were each considered as an inanimate, inhuman object attempting to replicate the appearance of a living human subject. Instead of viewing the faces as human heads that have been rendered in different ways, they were judged based on the different materials they represent, composed to create the form of human face. What material properties give the appearance of life and which appear inanimate? The basic question when viewing each of the images became: “Could this form have articulated movement, and could it move like a human given its material properties?” If the ability for movement is doubtful, the form becomes more of a sculptural representation of a face and less of an object trying to simulate real organic human tissue. Once the form is deemed sculptural, it is easier to accept as an abstract artistic representation.

The metal, glass, and porcelain faces easily fall into the category of inanimate sculptural objects with surface properties so different from human skin that they are unaffected by the uncanny valley. The meat and potato materials can be considered sculptural as well. However, they don't fall into this category as easily as the hard-surface man-made materials. These two surfaces are special cases: natural organic materials in an unnatural form. The meat shader is definitely unsettling for its similarity to flayed human flesh, but this comparison highlights a common aversion to blood and guts rather than the appearance of the uncanny valley. The juxtaposition of the realistic human eyes is truly what gives these materials any hint of life. Without the eyes, these faces would appear purely as formed lumps of their respective materials and nothing more. The eyes afford a second look and a bit of contemplation into what this creation actually is. In the end, these two faces appear as abstract inanimate portrait busts. They are admittedly created from rather odd materials, but they would not be considered sufficiently human-like to cause any confusion about their authenticity.

The remaining shaders are closer representations of real skin (wax, rubber, silicone, and video game) and are seen less as sculptural human forms and more as something attempting to be an actual human with capabilities for motion. These materials can be divided into surfaces that appear harder (rubber and video game) and softer (wax and silicone). The materials in both of these categories possess traits that could cause them to appear uncanny. The harder surfaces' lack of subsurface scattering and broad dull specular contribution makes them appear like hardened, dried-out skin with no blood underneath. The rubber shader gives the appearance of extremely thick makeup blocking any color on the skin and creating a very uniform surface reflectance. The video game shader is a bit further from real skin and looks more like a person wearing a foam latex mask. The more realistic eyes exacerbate this effect, making it appear like a different face might be behind this face, concealing its true nature. The soft materials, such as wax and silicone, have had any fine details and variation in the skin texture airbrushed away. The silicone surface represents pristine untouched skin that has never been affected by the elements in any way – a level of perfection and homogeneity that is not found in real human skin. The wax shader has a similar sterile, unpolluted surface quality. But in this case the subsurface scattering has been pushed even further than the already exaggerated silicone. Both materials have noticeably deeper light penetration, making the surfaces appear much softer than skin, but the wax allows light to penetrate so far that the face begins to appear as a highly translucent membrane like thick gelatin. The subsurface scattering, coupled with the more prominent specular highlights, help to give the surface the appearance of having a squishy feel.

For this investigator, the wax and silicone materials are the most disturbing. Silicone's overall shading quality can almost be mistaken for skin due to its similar light penetration and reflectance. It is no surprise that this synthetic material is widely used for making human dolls, skin prosthetics, humanoid robots, etc. However, the silicone skin in these examples vary substantially from human skin in regards to a lack of surface imperfections, lack of color variance, and lack of sub-surface coloration. The silicone material is very similar to skin in the way it interacts with light, but it has a texture and color that is very simplified and much too perfect to be real skin. This silicone effect is common in many films, such as the *Final Fantasy* series. The skin is represented as baby doll smooth with almost no color variation. This combined with an accurate skin-shading model creates an imbalance that causes the characters to appear uncanny.

This relates back to the argument for consistency of parts and level of detail, and why the photo-real face model created in this study is able to avoid the uncanny valley. One might reasonably argue that the tendency to want to airbrush out details and imperfections to make a face appear more beautiful is a cultural impulse. This lack of surface variation is more acceptable when the form of the face is stylized and unrealistic, but when the facial structure, the shading quality, and even hair, are at a photo-realistic level of detail, simplified skin texture stands out in stark contrast and is unsettling.

The disturbing quality of the wax material relates back to the question posed earlier: “Given the surface properties, is it conceivable that this character could have articulated movement?” The wax figure looks as if it could move, but its movement would be different from a human’s--perhaps like a person lacking a solid bone structure. The possibility that a character who is capable of movement would move in an unusual and possibly erratic way makes the idea of interaction very disconcerting. The unpredictable nature of the character, created by its surface quality, makes the face appear uncanny. Furthermore, the deep sub-surface allows light to penetrate far into the surface, but underneath the skin there is no indication of blood or musculature, or anything that would suggest there is a human armature. This only adds to the perception that the character is masquerading as a human, but is something else entirely.

6.2.2. The Eyes

Much like the skin, the eye shaders ranged in likeness from nearly photo-realistic to surreal. The different eye shaders were viewed side by side, all using the base skin shader as a constant for every render. The results found from analyzing the eye materials were surprisingly different from the skin.

When studying the various skin materials, surfaces like metal and glass that had properties very distant from real skin became sculptural in nature and were no longer viewed as an actual human. For this reason, they fell outside the uncanny valley since they did not appear to represent an actual living entity capable of movement and thought. In contrast, the eyes shaders that appear completely unrealistic could be arguably more disturbing than those closer to realism. As an overall assessment, the full spectrum of eye materials evoke a much more unsettling reaction than the skin materials. There is no single variation that could be regarded as falling completely outside the uncanny valley; they each have discrete properties that elicit a negative response.

The plastic eye material, which represents the variation closest to the realistic eye, is actually one of the most acceptable shaders in the spectrum. Nevertheless, its simplified details cause it to suffer from the same downfalls as the silicone skin shader. It has an unnaturally perfect surface, which is especially noticeable in the pristine eye whites. The glossy reflections and exaggerated refractions create a cute wide-eyed look that is more suitable for a stylized character.

The foggy glass material is not overly disturbing, but gives the odd appearance of eyes suffering from a medical condition. The foggy, glazed-over lens looks similar to an advanced form of cataracts, which will cause the eye to appear cloudy and inhibit vision. Cataracts are most commonly attributed to age, and are typically only seen in the elderly (Quillen, 1999). This representation of an aged, degrading eye, contrasts with the young face of the character and creates a strangely disquieting inconsistency, while evoking feelings of sickness, death and decay.

The gelatin shader doesn't appear as unsettling as some of the other materials. Its vibrant colors, exaggerated sub-surface scattering, and odd undulating specular pattern, give the surface a stylized look associated with non-photorealistic rendering, or “toon” shading. Although the actual gelatin material may appear similar to its real world representation, the exaggerated coloration may be too absurd and unrealistic to really upset the viewer.

The rubber and metal materials, which could both be considered rather unnatural based on their extreme deviation from a realistic eye, are effectively the most disturbing in the series of material variations. It is difficult to identify and state a singular reason why these particular eye representations elicit such a negative response; there are numerous possible contributing factors. The material properties of both shaders are clearly very different from realistic eyes. The rubber shader lacks any defined reflection and is overall very diffuse. The solid black pupils contrast against the overly simplified green and white color regions of the iris and sclera to create an unnaturally piercing, but at the same time detached or disconnected gaze. The color of the metal eye is strikingly different from that of the base eye. This is due to the fact that highly reflective metal receives most of its color from the environment it reflects. In this case, very little of the actual eye color is retained, resulting in a dark surface with stark highlights. The rubber material represents an eye that has too clearly defined diffuse color on the surface, whereas the metal material doesn't have any clear definition coming from the actual color of the eye. As a result, both eyes appear devoid of life. The metal eyes appear vacuous and ambiguous, lacking emotion and human thought. The rubber eyes look forever frozen in a solid gaze, cold and unnerving. Furthermore, the juxtaposition of unnatural eyes with realistic skin causes a disturbing uncertainty concerning the true form of the character. The rubber eyes appear as grotesque posthumous replacements for the original eyes of the character; an idea that suggests a ghastly narrative. The metal eyes imply a robotic mechanical form underneath. The eyes have been revealed to show an imposter wearing human skin.

7. CONCLUSION

Based on the results of this study, it seems that the uncanny valley can result from many contributing factors, and the range of its occurrence is somewhat nebulous. It appears that both theories on the uncanny valley do have some merit given the right conditions and the subjective perceptions of the viewer. It's difficult to claim whether the uncanny valley is the result of incorrect material properties, incongruous levels of detail and rendering quality, or simply disturbing subconscious connections to other images or ideas.

In the case of the skin shaders, it appeared that the material variations that were somewhat similar to realistic skin were more disturbing than the more abstract representations. This particular notion helps to validate the original theory on the uncanny valley made by Mori. The faces made of more implausible materials appeared to be mere sculptural representations of a human, and were rather nonthreatening based on their assumed inability to move or act in any way like a human. The materials similar to human skin, such as wax and silicone, produced the most uncanny effect and affect. These materials were disturbing based on their proximity to human skin. When considering silicone, the biggest problem was the inconsistent separation between surface texture and the shading quality. While the face interacted with light approximately the same way as skin, the surface of the face appeared far too pristine and without any imperfections. The wax material suffered from this same problem to a degree, but the most disconcerting aspect was its intense translucency, which suggested a non-human interior and an overly malleable surface. Both materials ultimately conveyed characters that were almost human, but given their surface properties, were assumed to have unpredictable motion and questionable cognitive abilities.

The eye shaders were more disturbing than the skin shaders as a group. There was no single variation that was completely free of uncanny qualities. Interestingly, the most unsettling eye variations appeared far from human-likeness. This was the complete opposite of the findings for the skin materials, and helps substantiate MacDorman's argument for a larger uncanny valley. Unlike the skin, it was more difficult to disregard the uncanny valley and claim that the more abstract eye variations were purely sculptural. The realistic skin, which was used as a constant, made the character believable as a human, but the unlikely eye materials created uncertainty and confusion towards its credibility. The most uncanny materials, metal and rubber, had greatly contrasting surface properties that were both very different from human eyes. The highly diffuse matte surface of rubber makes the eyes appear painted-on and unable to move based on the lack of a viscous surface. The metal lacks any color definition from the surface, and is an ambiguous reflection of the environment. The unnerving dry stare of the rubber eyes and the vacuous mechanical image of the metal eyes evoke thoughts of death. Beyond the material properties, these eye representations—together with human skin—create unsettling perceptions. The metal draws the instant impression of a robotic form under the surface, while the rubber appears as a crude replacement for real eyes.

The results of the study show me that certain material properties are not inherently more conducive to triggering the uncanny valley effect, but that the phenomenon is more the result of the tangible and intangible associations these materials create, as well as an inconsistency in detail. Materials properties that influence the assumed motion of the character in a negative way, or disconcert how the surface would feel to the touch, can account for an uncanny response. A dissonance between the shading quality of a surface and its texture detail can also have an adverse effect. A photo-real shading model leads to the expectation of high fidelity textures and subtle surface shifts that inform blemishes, pores, discoloration, etc. This applies to the different parts of the entire character and not just a single surface. Also, when any given aspect of a CG model is recognizably accomplished, or lacking in quality, the combined image breaks down and is unsuccessful as a whole. All aspects of the face should be at a consistent level of detail to ensure that nothing can be singled out and analyzed. Finally, the uncanny valley seems a result of the perceptions of the viewer and what correlations they derive based on the images. Certain surface combinations evoke more negative feelings of death and abhorrence based on what they represent in the viewer's mind, as well as the story that is created from the images.

8. FUTURE WORK

The number of factors that can affect realism in a CG face and influence the uncanny valley is undoubtedly larger than the scope of this study. Only the shading and texture of the skin and eyes were explored. Aspects of the face such as form, lighting, and especially motion could benefit from a deeper exploration of their relation to the uncanny valley.

A direct extension of this thesis would be to quantitatively study the different face material variations. Numerous participants could be presented with the images to determine different perceptions of what is considered uncanny. These varied perspectives could allow for a more substantiated conclusion regarding the occurrence of the uncanny valley.

The facial model remained geometrically constant throughout the exploration; there were no changes made whatsoever to its form. Many past studies have thoroughly analyzed the overall shape of the face by changing the relative size and position of prominent features and studying viewer reaction. This earlier research examined the face on a macro scale, inquiring only how more noticeable shifts in geometry affect the uncanny valley. In the current generation of visual effects and games, creating a high quality model of the face is not as difficult as it once was. The proliferation of 3D scanning technology, which has been widely adopted by those who have an interest in replicating reality, can allow for a nearly perfect recreation of the human face with almost no human interaction. However, geometry generated from these scans has messy topology that requires heavy modification to be usable in a production environment. During this process of transferring the scan information to more uniform geometry, some of the subtle details can be lost, and it is the artist's job to fix these issues. Areas such as the eyelids, lips, ears, nostrils, and high frequency detail such as pores and blemishes are all delicate features that require an artist's attention to achieve realism. Studying how these micro details affect the uncanny valley can still be very relevant to the current state of the CG industry.

From an artistic point of view, incorrect lighting can completely destroy the believability of a face, regardless of accuracy in shading and geometry. Research exploring the face in different lighting environments and with different types of lights and light properties could greatly benefit facial realism. Studying how lighting can alter a viewer's emotional response, or how it can change their comfort level with a face can help further an understanding of the uncanny valley.

Perhaps the most important contribution to understanding the uncanny valley in CG characters is realistic motion, a subject that was not explored in this thesis. When studying the skin, one of the biggest concerns was the assumed motion of the character with its given material properties. The fact that an inanimate character can evoke fears of motion proves how important correct movement is to creating a character with which viewers can feel comfortable. In this study, many skin variations were perceived as sculptural since they appeared to not have the ability to move. However, if these forms were animated, proving that they could in fact move, the result could be much more uncanny. The form creates expectations of movement, and if these expectations do not match the reality, the result could be a further rejection by the viewer (Mori, 1970). Studying the different skin and eye variations on a character in motion could be a way to further advance this specific type of research.

Motion of the face could benefit from exploration of details at a micro level. The movement of the eye, as well as the small motions of the face can be extremely subtle, but their absence or inaccuracies can destroy believability. Looking at the same motions with slightly different timings or levels of exaggeration on a realistic character could be another worthwhile exploration. Also, studying secondary motion such as the elasticity and folding of the skin, and how it moves correctly across the underlying bone structure is required to further advance current facial animation. These motions are not conspicuous, but are likely very important in allowing a CG character to finally bridge the uncanny valley.

REFERENCES

Birn, J. (2000). Digital Lighting and Rendering. IN: New Riders Publishing.

Dargis, M. (2011, December 20). Intrepid boy on the trail of mysteries. The New York Times, p. C1.

Dohlman, C. A. (1971). The function of the corneal epithelium in health and disease. Investigative Ophthalmology, 10(6), 383-407.

Freud, S. (1955). The "uncanny". London: Hogarth Press / Institute of Psychoanalysis. (Original work published in 1919).

Green, R. D., MacDorman, K. F., Ho, C., & Vasudevan, S. (2008). Sensitivity to the proportions of faces that vary in human likeness. Computers in Human Behavior, 24(5), 2456-2474.

Igarashi, T., Nishino, K., & Nayar, S. K. (2005). The appearance of human skin. Tech. Rep. CUCS-024-05, Department of Computer Science, Columbia University, June.

Jentsch, E. (1906). On the psychology of the uncanny. Angelaki: Journal of the Theoretical Humanities, 2(1), 7-16.

Kaba, F. (2013). Hyper-realistic characters and the existence of the uncanny valley in animated films. International Review of Social Sciences and Humanities, 4(2), 188-195.

Li, H. H. (1980). Refractive index of silicone and geranium and its wavelength and temperature derivatives. Phys. Chem. Ref. Data, 9(3), 561-658.

Lidwell, W., Holden, K., & Butler, J. (2003). Universal Principles of Design. MA: Rockport Publishers.

MacDorman, K. (2006). Subjective ratings of robot video clips for human-likeness, familiarity, and eeriness: An exploration of the uncanny valley. Proceedings of the ICCS/CogSci-2006 Long Symposium: Toward Social Mechanisms of Android Science. Vancouver, Canada.

MacDorman, K. F., Green, R. D., Ho, C., & Koch, C. T. (2009). Too Real for Comfort? Uncanny Responses to Computer Generated Faces. Computers in Human Behavior, 25(3), 695-710.

Malitson, I. H. (1965). Interspecimen Comparison of the Refractive Index of Fused Silica. JOSA, 55(10), 1205-1208.

Mori, M. (1970). Uncanny valley. Energy, 7(4), 33-35.

O'Neill, R. (2008) Digital Character Development: Theory and Practice. MA: Morgan Kaufmann Publishers.

Patel S., Marshall J., & Fitzke F. W. (1995). Refractive index of the human corneal epithelium and stroma. J. Refract. Surg. 11(2), 100–105 (1995).

Pollick, F. E. (2010). In search of the uncanny valley. Lecture Notes of the Institute for Computer Sciences. Social Informatics and Telecommunications Engineering, 40(4), 69–78.

Quillen, D. A. (1999) Common causes of vision loss in elderly patients. Am Fam Physician, 60(1), 99-108.

Ralph, C., & Chau, I. (1993). Mullite. Retrieved from http://www.mindat.org/min-2806.html.

Rhodes, G., & Tremewan, T. (1996). Averageness, exaggeration, and facial attractiveness. Psychological Science, 7(2), 105-110.

Seyama, J., & Nagayama, R. S. (2007). The uncanny valley: The effect of realism on the impression of artificial human faces. Presence: Teleoperators and Virtual Environments, 16(4), 337-351.

Sugimoto, K., & Matsuda, S. (1983). J. Electrochem. Soc., 130, 2323.

Travers, P. (2001). Final Fantasy. Rolling Stone. Retrieved from http://www.rollingstone.com/movies/reviews/final-fantasy-20010706.

Watanabe, S., Kakigi, R., Koyama, S., & Kirino, E. (1999). It takes longer to recognize the eyes than the whole face in humans. NeuroReport, 10(10), 2193-2198.

Wilhelm, K., Elsner, P., Berardesca, E., & Maibach, H. I. (1997). Bioengineering of the Skin: Skin Surface Imaging and Analysis. FL:CRC Press.

Wood, L. A., & Tilton, L. W. (1949). Refractive index of natural rubber for different wavelengths. National Bureau of Standards, 43(1), 57-64.