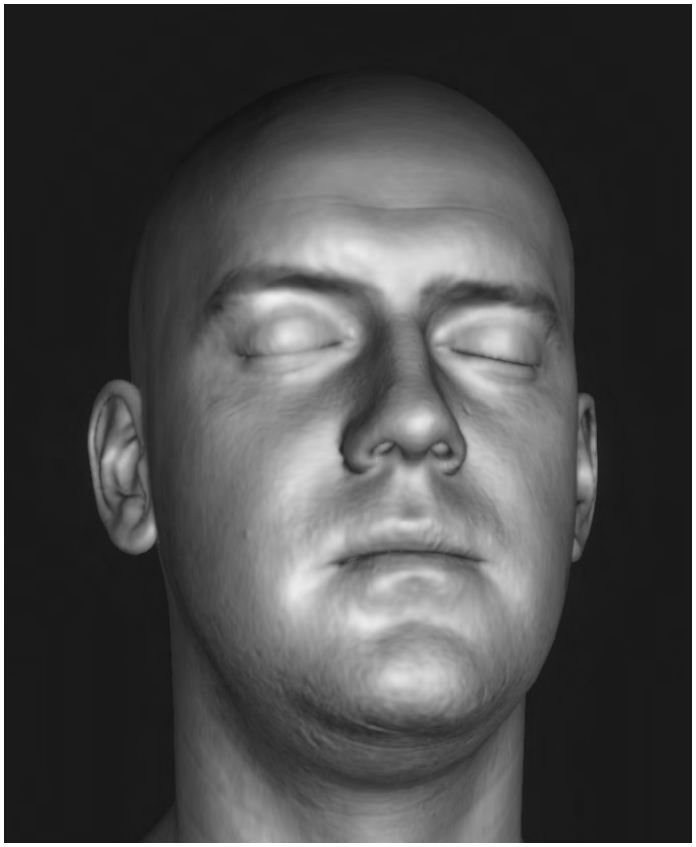

Simulating subsurface scattering in skin by near-infrared light

Master's thesis in Computer science and engineering

DANIEL FELCZAK

Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
UNIVERSITY OF GOTHENBURG
Gothenburg, Sweden 2020

Master's thesis 2020

Simulating subsurface scattering in skin by near-infrared light

DANIEL FELCZAK

UNIVERSITY OF
GOTHENBURG

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Computer Science and Engineering
Chalmers University of Technology
University of Gothenburg
Gothenburg, Sweden 2020

Simulating subsurface scattering in skin by near-infrared light
DANIEL FELCZAK

Supervisor: Erik Sintorn, Department of Computer Science and Engineering
Advisor: Per Svensson, Smart Eye
Examiner: Ulf Assarsson, Department of Computer Science and Engineering

Master's Thesis 2020
Department of Computer Science and Engineering
Chalmers University of Technology and University of Gothenburg
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Real time render with subsurface scattering behavior of human skin according to real-life measurements.

Typeset in LaTeX
Gothenburg, Sweden 2020

Simulating subsurface scattering in skin by near-infrared light
DANIEL FELCZAK
Department of Computer Science and Engineering
Chalmers University of Technology and University of Gothenburg

Abstract

The reflectance properties of human skin in the visible spectrum and near-infrared spectrum are significant. In general, skin appears more uniform. In this thesis we have chosen to explore the difference between subsurface scattering in near-infrared and visible light and also how to implement subsurface scattering in a real-time setting. We measure a diffusion profile of Caucasian skin in near-infrared light at 940 nm on the inside of the forearm. This diffusion profile is used with existing real-time skin rendering algorithms to implement subsurface scattering. Our diffusion profile of 940 nm did not show a major difference to that of red light (700 nm), however, it is not conclusive as measurements of red light were not our own. Further, when implementing subsurface scattering we used previous research and do not introduce any novel implementation. Finally, we conclude that a large part of the increased uniformity can be attributed to the fact that the skin is only being illuminated by a single wavelength (940 nm). Consequently, small details, normally highlighted by absorption in the visible spectrum become far less noticeable. However, near-infrared light also penetrates deeper into the skin, highlighting other details, e.g., blood vessels in certain areas of the face.

Keywords: near-infrared, subsurface scattering, computer graphics, real-time rendering, diffusion profile.

Acknowledgements

First I would like to thank the people at Smarteye who allowed me to have a very welcomed stay during the period I worked in their office. I would also like to thank my supervisor Per Svensson at Smarteye for being helpful, providing feedback and discussion throughout the entire thesis. Then I would like to thank Jörgen Thaung at Smarteye for helping us conduct measurements to be used in the thesis. Finally, I would like to thank Erik Sintorn, my supervisor at Chalmers, for his continued guidance and feedback on the thesis work.

Daniel Felczak, Gothenburg, June 2020

Contents

List of Figures

List of Tables

1
Introduction

1.1 Background

Since the vast increase of computational power which has in the recent two decades exponentially increased, the demand for increased accuracy in simulated environments have in a similar fashion followed. This is obvious if we look at the ubiquitous use of computer-generated effects in the movie industry but also in the video-game industry. Both of these industries are pushing the quality of realistic effects and environments in order to bridge the gap between reality and simulation. Moreover, this demand is not exclusive to the various entertainment industries but is also important for many different areas of research, e.g., allowing self-driving vehicles to learn in realistic simulated environments and other cases where, in general, it is beneficial to reduce the gap between reality and simulation in order to allow algorithms to operate in a realistic environment.

Today, a majority of rendering accomplished in the computer graphics community predominantly use mathematical models and data which describe the interaction between light and materials in the *visible spectrum*, i.e., the wavelength range of electromagnetic radiation which we humans can interpret with our eyes. However, in the field of computer graphics, the domain outside visible light remains relatively unexplored. In this thesis, we explore the interaction between human skin and light which exists outside the visible spectrum, more specifically, we explore the interaction between light and human skin where the skin is illuminated by light in the *near-infrared* wavelength range. Due to human skin being a translucent material, the major topic in this thesis will the subsurface scattering phenomenon between near-infrared light and human skin.

1.2 Motivation

This thesis is motivated by a real-time application for a driver-monitoring system aimed at the commercial vehicle industry. This system monitors a person's current physical and mental state, e.g., unresponsiveness, tiredness etc. This is achieved through face tracking and various machine learning algorithms. For the software to draw conclusions about the state of the person, it requires a consistent stream of images where the face is illuminated such that it can capture the important tracking features of the face. In order for the illumination not to be distracting or cause annoyance to the person being illuminated, the system utilizes near-infrared

light. In addition to this software, the company developed a real-time 3D software application with the purpose of acting as a flexible tool for testing and generating realistic data. Currently, the application does not implement any functionality which simulates the scattering effects between near-infrared light and human skin. Instead, the application currently only apply light models which are typically used in the visible spectrum. In order for the algorithms to operate in an increasingly accurate, or realistic, environment, the vision is to include functionality which simulates the behavior of near-infrared light and human skin.

1.3 Problem Statement and Goal

The resulting problem statement for this thesis is to investigate how near-infrared light interacts with skin and attempt to measure the scattering properties of near-infrared light as it interacts with human skin. According to D'Eon et al. [4] demonstrate that as the wavelength increases in visible light, so does the degree of subsurface scattering, additionally, Cooksey et al. [3] show that light in the near-infrared range scatters further into the skin than that of the light in the visible spectrum. Thus, for this thesis, the hypothesis is that near-infrared light scatters further into the skin and potentially causes a larger subsurface scattering effect than light in the visible spectrum. Moreover, observations between real-life images of a face illuminated by near-infrared and visible light, separately, indicate that certain areas likely produce a larger degree of subsurface scattering. However, since data regarding this is either hard to gather from previous research in other fields, e.g, optics, or is very scarce in general, the goal is to conduct measurements and gather data about the subsurface interaction between human skin and near-infrared light. The gathered data will be applied in a real-time application with the goal of increasing the accuracy and attempt to capture the scattering effect of near-infrared illumination with human skin.

1.4 Limitations

For this thesis we will limit ourselves to exclusively measuring the subsurface scattering effect between human skin and near-infrared light. There are multiple other details which are different between visible light and near-infrared light and their respective interactions with skin, as will be discussed in more detail in the theory section.

Focusing on the subsurface scattering component is motivated largely by a reasonable assumption that the increased penetration depth by near-infrared light increases the degree of subsurface scattering. Additionally, in the current state of the software application, it does not incorporate any sort of approximation of the subsurface scattering effect, hence, it seemed reasonable to implement in order to increase realism.

Another limitation is that the measurements will be conducted at a wavelength

of 940 nm. This is what is used in the company's product and is readily available. Further, due to the restrictiveness of the measurement setup, only the measurements on the inside of the forearm are used. This may not accurately represent the scattering behavior of the face, but the method can, if one desires, be extended for other parts of the body.

1.5 Outline

In Section 2, we will briefly go over previous work both in conducting measurements for subsurface scattering purposes. Additionally, the previous work within real-time subsurface scattering is reviewed. In Section 3 will cover the theory which this thesis will touch upon, e.g., components of scattering, subsurface scattering and diffusion profiles. Further, we mention briefly the properties of human skin and some general details of near-infrared light. Section 4 covers how real-life measurements between human skin and near-infrared light was conducted, then we show how this was implemented as a shader in Unity. In Section 5, the results are presented. Finally we conclude with discussion and conclusion in Section 6.

2
Previous Work

We review previous work regarding the spectroscopy of human skin, how to measure subsurface scattering of materials, and methods to implement subsurface scattering for real time rendering.

2.1 Near-infrared light and skin

Near-infrared light has been used in many different type of areas, both in engineering and research. In the field of computer vision it has been used as illumination for face detection hardware [15]. In optics, there has also been research to detect skin in images, by utilizing the unique reflectance properties of human skin in the visible and near-infrared wavelength spectrum, it is then possible to train use neural networks to identify human skin in images [18][14].

The behavior between near-infrared and visible light is quite significant. Kansawa et al. [14] report that the reflectance at lower wavelengths (350 nm) is significantly lower than that in higher wavelengths (1090 nm). In addition, Cooksey et al. [3] report similar results, but they also report that the penetration depth increases at longer wavelengths reaching a maximum penetration depth of 3.5 mm at a wavelength of 1090 nm. Mohamad et al. [19] conducted measurements of the moisture contents of skin, they show that at higher wavelengths near-infrared light will have peaks where it will be absorbed in the water contents of the skin. These absorption peaks happen at 970 nm, 1100 nm, 1450 nm, and 1900 nm. The absorption peaks do not vary much between different areas of the body.

Igarashi et al. [8] give a thorough review of human skin and the current practices for rendering human skin. Anderson et al. [1] show from experiments, that most of the scattering occurs in the dermis of the skin, and how the scattering depth increases with longer wavelengths.

2.2 Measurement of materials and near-infrared light

The focus in this thesis is to capture the subsurface scattering in human skin. A common method to separate the subsurface scattering component is presented by Jensen et al. [10]. Their method captures the subsurface scattering properties of

human skin, but could also be applied to other translucent materials where light is more likely to scatter than be absorbed. They are able to generate so called *diffusion profiles* which describe, as a function of distance from the initial point of entry, how energy dissipates as it scatters in the translucent material. In a similar effort, Tariq et al. [24] also measure the scattering profile of different areas of a human face. Their approach allows them to capture the scattering profile of the material but their measuring setup is more complex to adapt. Both of these methods build on the assumption that the material exhibits the property of having significantly larger amounts of scattering than absorption. Weyrich et al. [25] also provide extensive detail on how they generated realistic skin. In order to capture the translucency and scattering behavior of skin, they constructed their own measurement device consisting of linear optical fiber sensors. After the measurements they generate a diffusion profile with the theoretical foundations provided in the paper by Jensen et al. [10].

The interaction of near-infrared light and various materials is presented by Choe et al. [2]. They demonstrate the different and unique phenomenon, relative to visible light, that occurs when near-infrared light interacts with materials. In many cases the materials lose much of their texture. A dark fabric of a cotton shirt with various printed logos, for instance, will appear entirely bright and uniform. They also present the measurement setup using near-infrared light as a light source for data gathering of a large number of different materials, including human skin. However, their measuring method generates a BRDF of the given material but does not separate the subsurface scattering component.

2.3 Real-time subsurface scattering

For real-time implementations of subsurface scattering, recent research use the diffusion approximation as presented by Jensen et al. [10] which builds on the assumption that if the translucent material is optically dense such that multiple scattering dominates, the material will tend toward uniformity in its scattering distribution. D'Eon et al. [4] present a real-time implementation where they utilize diffusion profiles. They fit the profile data with a sum-of-Gaussians approximation. Then they apply each Gaussian as filter in texture-space. The end result is essentially an averaging filter controlled by the diffusion profile. In a similar fashion, Jimenez et al. [11] build on the same idea of using diffusion profiles but their approach operates in screen-space instead of texture-space. The same authors introduce a method called separable subsurface scattering. In this method they extend their screen space approach and uses the fact that in a sum-of-Gaussian function fit, some Gaussians contribute significantly more than others, consequently, their method essentially excludes the negligible Gaussians, thus making it more efficient [13] as the number of filtering passes are reduced.

Penner et al. [21] noted that subsurface scattering is more pronounced in areas where the surface contains some degree of curvature and non-uniform lighting. In other more flat and uniformly illuminated areas, the surface behaves similar to that

of a diffuse surface. Consequently, they generate a lookup texture using the diffusion profile for the target material by computing the falloff of incoming light on circles with varying radii. When shading a pixel on a mesh with arbitrary curvature, they fit a sphere to the curvature and use the lookup texture based on the radius of the sphere to determine the degree of subsurface scattering.

3
Theory

In the following sections, the relevant topics for this thesis are explained. The different types of scattering events are discussed with a focus on the phenomenon of subsurface scattering. The scattering events discussed here are not an exhaustive list. For more details we refer to *Physically Based Rendering* by Pharr et al. [22]. Further, we discuss the details of human skin in the context of computer graphics. A more complete review on the topic of skin and rendering of human skin can be found in *The Appearance of Human Skin: A Survey* by Igarashi et al. [8]. Finally, we discuss the interaction between near-infrared light and human skin.

3.1 Scattering events

For real-time purposes, the different types of interactions between material surfaces and light are commonly separated into three different components. *Diffuse scattering*, which is when light scatters in random directions due to a surface roughness. This roughness causes the outgoing rays to be perturbed in random directions. This is also known as *Lambertian* scattering. The other type is *specular scattering*, which is the event when light interact with smooth surfaces. In this case, we can consider it the complete opposite of diffuse scattering as light will instead exit in a specific direction. This new direction is easily found if the position of the light source is known as well as the angle of incidence to the surface. If the surface is perfectly smooth (e.g., a mirror) then the exiting angle will be exactly equal to that of the incoming angle, relative to the surface. The last type of scattering event mentioned here is that of *transmission*. Transmission is when an incoming ray of light arrives at a surface and transmits through, or into, the material. Depending on the material and refraction index, it will either be slightly perturbed from its original incoming direction, or traverse through the material in its original path. Once transmitted, the next action depends on the properties of the current medium. For instance, when light travels in a medium like water, the ray of light is likely to continue to travel in its perturbed path until it exits the water or become absorbed along the path. This is commonly referred to as *volumetric scattering*. Subsurface scattering on the other hand is also an event of transmission in which light transmits into a medium. However, due to the material properties (e.g., skin contains fibers and a multitude of particles which light can interact with), it will be much more likely to continue to scatter until it exits the medium rather than being absorbed. Theoretically there is no difference between subsurface scattering (e.g., in human skin) and volumetric scattering (e.g., in water). They both use the same theoretical framework

in computer graphics. The distinction between the two events is based on the type of material, or medium, light is interacting with. With subsurface scattering, the incoming light ray will scatter inside the material and eventually exit close to the original point of entry, whereas in water, light is more likely to continue to travel relatively far without a large degree of perturbation from its initial direction until it either exits or becomes absorbed.

3.1.1 Diffuse scattering

Diffuse scattering occurs when light arrives at a surface and is approximately equally likely to exit in any direction due to the micro-roughness of the material. In this case there is no form of transmittance. Instead, at the event of scattering, the new direction of the ray will be approximately random in the hemisphere above the point of intersection with the surface, see Figure 3.1.

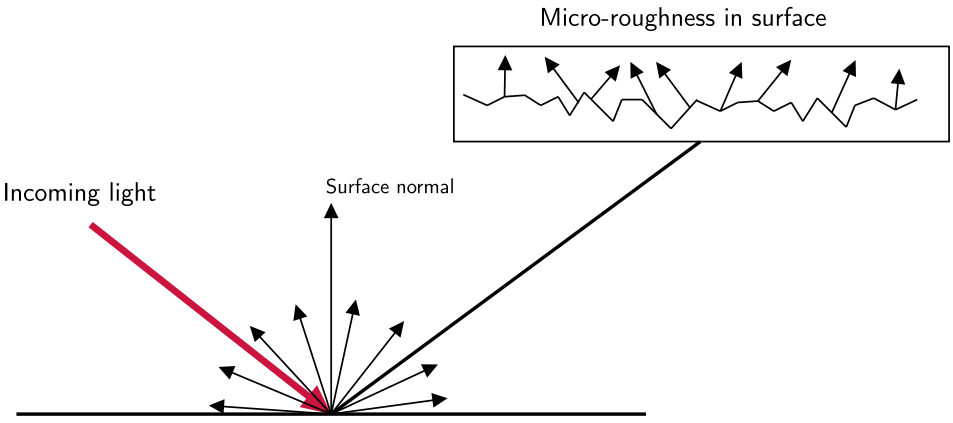

Figure 3.1: Diffuse scattering, the incoming ray has an equal probability of exiting in any direction around the hemisphere of the original point of intersection.

3.1.2 Specular scattering

The second type of scattering is specular scattering. In this case we deal with surfaces which are smooth such that many of the rays will travel in a similar outgoing direction. A common example of such a material is that of a mirror. All interactions of light with a perfectly smooth mirror will travel bounce exactly according to the law of reflection (i.e., outgoing angle of direction equals the incoming angle of direction, relative to the surface normal). Most materials in the real world exhibit some degree of specularity (e.g., the top layer of human skin), however, there is often a level of micro-roughness which will perturb the light from its mirror reflection as seen in Figure 3.2. In this case, light scatters in a lobe around the perfect reflection direction. Materials and surfaces which exhibit this sort of behavior are commonly referred to as glossy surfaces.

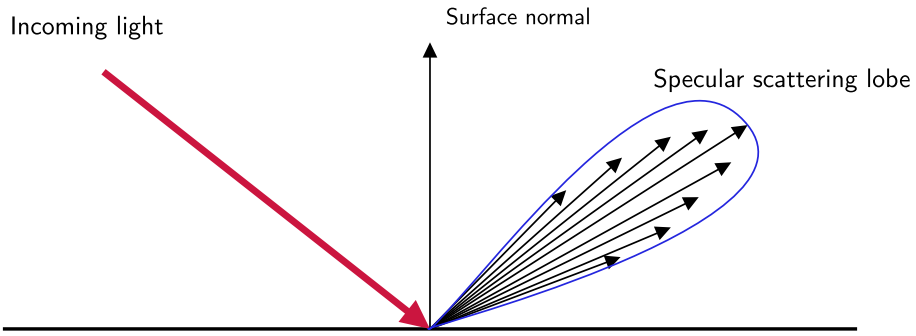

Figure 3.2: Specular scattering: when light interacts with smooth surfaces it will exit the surface in an angle around the perfect reflection direction. Real materials will however often scatter in a lobe around the perfect reflection direction due to small perturbation in the surface.

3.2 Subsurface scattering

Now we consider the phenomenon of subsurface scattering. This event occurs for translucent materials and specifically for those which are optically thick. Inside an optically thick material, light has a much higher probability of continuing to scatter rather than being absorbed. Even though the underlying theory is equivalent to the more general phenomenon of volumetric scattering, subsurface scattering only concerns materials where light scatters close to the surface of the medium. After a large number of scattering events, a ray of light, will eventually exit near the initial point of entry. Some examples of medium which exhibit this property are human skin, potatoes, milk and marble. Subsurface scattering can be further divided into two subevents: single scattering and multiple scattering.

3.2.1 Single and Multiple scattering

Single subsurface scattering is defined as light which scatters only once inside the media before it exits again, see Figure 3.3. Multiple subsurface scattering is defined as light which scatters multiple times inside the media before exiting, see Figure 3.4. Skin is a complex multi-layered material where each layer contains multiple sublayers. Inside each layer of the skin exists various cells, different fibers, blood, water, melanin and more, which light will interact with as it travels inside the skin. It should also be noted that the number of scattering events depends largely on the wavelength of the light. For instance, blue light is more likely to be absorbed in the water in our skin than light in the red wavelengths. The smoothing effect which is often associated with subsurface scattering, is the result of multiple scattering. This is due to the large amounts of scattering inside the medium which spread the incoming rays of light into a small area around the initial point of entry.

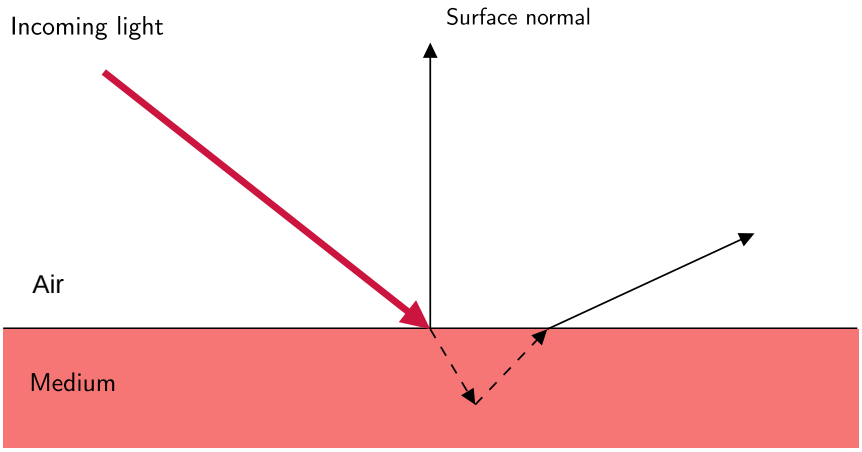

Figure 3.3: Single scattering event, after transmitting from air into the medium, light scatters once inside the medium before exiting at a different location from the initial point of intersection.

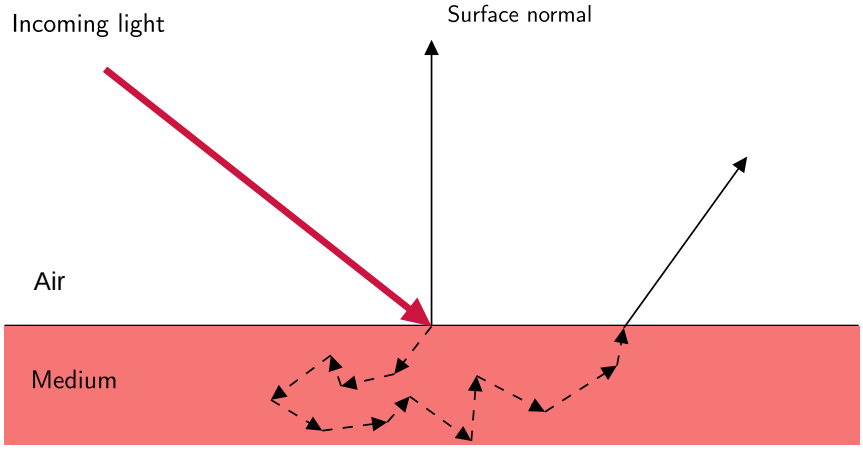

Figure 3.4: Multiple subsurface scattering. The incoming ray of light scatters at the surface, transmits through the material, then continues scatter inside of the medium until it exits. Note that events of absorption can occur during this path of scattering but this is not depicted in the figure.

3.2.2 Absorption

Absorption is the phenomenon when the light travels inside the material and does not continue to scatter but is instead absorbed. When the ray of light is absorbed it is converted to some other form of energy, e.g., thermal energy. Given a material, its spectrum of absorption and the type of incoming light will define the appearance of the material. For instance, blood, illuminated by the visible spectrum of light, in air, is a medium which will absorb light in the blue wavelengths (400 nm). This is what results in the red appearance of blood. However, if we now instead assume that blood is observed in deep water, where much of the red light is absorbed, then it will instead have a green color. This is due to the fact that water absorbs a large degree of red light. This is what gives water its blue appearance. Thus, regardless of the spectrum of absorption of the material, the incoming spectrum of light is also a defining factor for the appearance of the material.

3.2.3 Computational complexity of subsurface scattering

In order to compute the scattering at a given point on a surface, the model which incorporates subsurface scattering is the *bidirectional subsurface scattering reflectance distribution function* (BSSRDF). This is defined as:

$$S(x_o, \omega_o, x_i, \omega_i) = \frac{dL_o(x_o, \omega_o)}{d\Phi(x_i, \omega_i)} \tag{3.1}$$

This function describes the ratio of how much energy exits, given an incoming amount of energy. x_i is the incoming ray of light and ω_i is the incoming direction of the ray. x_o is the outgoing ray of light and similarly, ω_o is the outgoing direction of the ray. As we can see, the function has many degrees of freedom and further, if one wishes to compute the energy exiting from a given point in space in a given direction, this requires the evaluation of the following double integral:

$$L_o(x_o, \omega_o) = \int_A \int_\Omega S(x_o, \omega_o, x_i, \omega_i) L_i(x_i, \omega_i)(n \cdot \omega_i) d\omega_i dA \tag{3.2}$$

this expression tells us to arrive at the energy exiting from x_o in direction ω_o, we need to integrate the surface of the entire object A, over the hemisphere at x_i, Ω. The BSSRDF is used here as a factor of the incoming light to compute how much incoming energy $L_i(x_i, \omega_i)$ will exit. $(n \cdot \omega_i)$ is known as the cosine attenuation factor or Lambert's cosine law. $L(x_i, \omega_i)$ is the incoming radiance arriving at x_i from direction ω_i.

To evaluate this integral, the typical approach is to use Monte Carlo integration which requires a large set of samples before converging to the true average. This consequeuntly leads to considerable computational times and it is not suited for real-time purposes. Further, even in a offline rendering process, it is still a relatively expensive phenomenon to compute. The use of diffusion profiles as a method to approximate the degree of subsurface scattering considerably decreases the computational requirements with minimal visual error. However, it should be noted that although it decreases the rendering times, the required time for each frame is still

too long to be considered for real-time purposes. The approach for computing the subsurface scattering component in a real-time setting is done with approximative methods combined with the data of diffusion profiles. Diffusion profiles build on the assumption that, for materials that are highly scattering, the scattering will tend to become isotropic. After a large number of scattering events, the initial point of entry becomes a negligible factor in future scattering events. Finally, when the light exits the medium, its scattering behavior is considered similar to that of diffuse scattering. Further, this approximation is only applicable for materials where the scattering coefficient is comparatively much larger than the absorption coefficient.

3.2.4 Diffusion profiles

Diffusion profiles are a concept introduced by Jensen et al. [10]. They were motivated to reduce the rendering time of subsurface scattering models for translucent materials. As mentioned in the previous section, in order to render translucent materials with physical accuracy, the computational effort is large. This is due to the number of scattering events occuring. For instance, in human skin, it is upwards of a hundred events before exiting [10]. Jensen et al. introduced a method where the subsurface scattering effect can be approximated with a *dipole approximation* of the BSSRDF. Their method relies on separating the single scattering and multiple scattering effect into separate components. Note that we will focus on the multiple scattering part here. The multiple scattering is what contributes to the smoothness of the material. More importantly, it is also the most computationally intensive part as it requires many bounces inside the medium. The multiple scattering part of the BSSRDF, S_d is approximated by the following:

$$S_d(x_o, \omega_o, x_i, \omega_i) = \frac{1}{\pi} F_t(\eta, \omega_i) R_d(\|x_i - x_o\|) F_t(\eta, \omega_o) \tag{3.3}$$

where S_d is the multiple scattering term of the BSSRDF, x_i is the initial point of entry, x_o is the point of exit after light has scattered multiple times in the skin. R_d is the diffusion profile, F_t are the Fresnel terms which are multiplied to account for transmission in and out of the medium. The *diffuse subsurface reflectance* term R_d is defined in the following manner:

$$R_d(r) = \frac{\alpha}{4\pi}\left(z_r\left(\sigma_{tr} + \frac{1}{d_r}\right)\frac{e^{-\sigma_{tr} d_r}}{d_r^2} + z_v\left(\sigma_{tr} + \frac{1}{d_v}\right)\frac{e^{-\sigma_{tr} d_v}}{d_v^2}\right) \tag{3.4}$$

The details of this equation can be found in the original paper by Jensen et al. [10] and Mukaigawa et al. [20] provide more succinct details on the equation. In short, the free parameters of the equation are the scattering coefficient σ_s and the absorption coefficient σ_a. In addition, there is an η, the refractive index, this parameter must be defined empirically from measurements.

The single scattering component and the multiple scattering component are evaluated separately but both rely on using the diffusion profile R_d. The profile is used to sample an outgoing point x_o given a point of entry x_i. To sample the multiple scattering component, they use typical Monte Carlo integration [22], where, given

that we have some initial point of intersection x_i, the samples are chosen from the density distribution,

$$\sigma_{tr} e^{-\sigma_{tr} d} \tag{3.5}$$

yielding a sample x_o at some distance d away from x_i. The Equation 3.5 contains the *effective extinction coefficient*, σ_{tr}. The coefficient σ_{tr} is dependent on the absorption coefficient σ_a and the scattering coefficient σ_s. Both of these coefficients are generated from measurements and define the behavior of the density. The single scattering sampling details can be found in the original paper [10].

3.2.5 Measuring diffusion profiles

If the scattering and absorption coefficients are known, it is possible to compute a diffusion profile using Equation (3.4). This equation describes how energy will dissipate as light scatters further from the initial point of entry. A diffusion profile can also be measured directly with a method introduced in the same paper by Jensen et al. [10]. They focus part of a light source to a small point such that they are able to separate the single scattering effect with the multiple scattering. An example of this is shown in Figure 3.5.

Figure 3.5: Human skin being illuminated by a focused point of near-infrared light.

Using Equation (3.4), we can generate the following diffusion profiles for human skin in red, green and blue light. After performing our own measurement we will use this as a reference in order to verify our measurements.

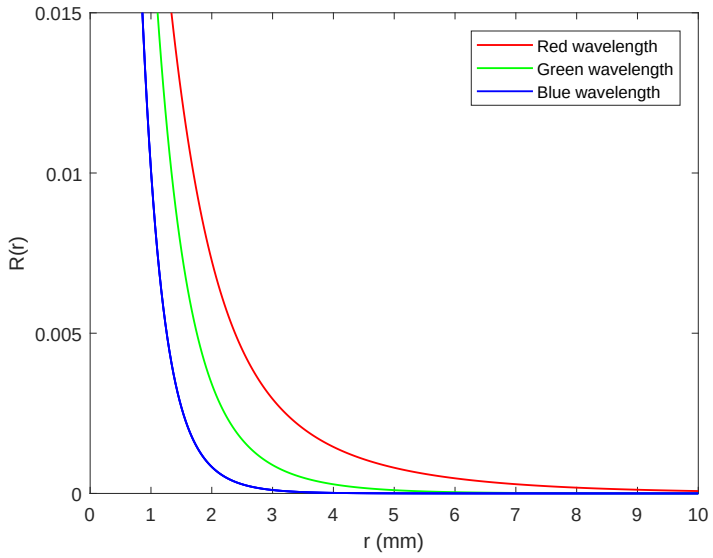

Figure 3.6: Diffusion profile for skin at different wavelengths with $\eta = 1.3$ and following scattering and absorption coefficients: **Red**: $\sigma_s = 0.74$, $\sigma_a = 0.032$, **Green**: $\sigma_s = 0.88$, $\sigma_a = 0.17$, **Blue**: $\sigma_s = 1.01$, $\sigma_a = 0.48$

3.3 Subsurface scattering in human skin

In the following section we limit ourselves to discussing the details of skin which we deem are most relevant to rendering. Skin is typically defined to contain three large layers with multiple sublayers. In the context of rendering, however, it is common to use a three-layer model [6]. The obvious features of the top layer are pores, hair, moles and wrinkles. There is also a layer of oils and lipids which contribute to the specularity of skin. Further, we have the mid-layer, the *epidermis*, and the bottom layer, the *dermis*. The epidermis and the dermis both consist of multiple sublayers. Thorough details of the sublayers contained in the epidermis and dermis can be found in the work by Igarashi et al. [8].

3.3.1 Epidermis

The epidermis is a relatively thin layer. The thickness will vary depending on which part of the body is measured, but on average it is measured to be 0.2 mm. In the epidermis a large component which affects the appearance of skin is that of melanin which only exists in the epidermis. Melanin is a *chromophore* which is defined as a light-absorbing chemical compound. The skin color of humans is highly dependent on the relative amount of melanin in the epidermis. Melanin is thus one of the main absorbers in human skin. The amoount of absorption decreases as the wavelength increases. Specifically, beyond 1100 nm, the melanin absorption abilities becomes almost negligible between different types of skin colors [1]. Besides melanin, there are also keratin fibers which hold water that also contribute to the absorption in the epidermis.

3.3.2 Dermis

Compared to the epidermis, the dermis is a much thicker layer. It varies in thickness between 1-4 mm on different areas on the human body. The main components in the dermis are collagen fibers. These fibers are the primary contributors to the fact that skin is a highly scattering material [1]. In addition, the dermis contains a large network of capillaries and veins which carry hemoglobin. Hemoglobin is another chromophore that accounts for absorption in skin. Finally, the dermis also contains water which contributes to absorption in the visible and near-infrared wavelength range. [8][18].

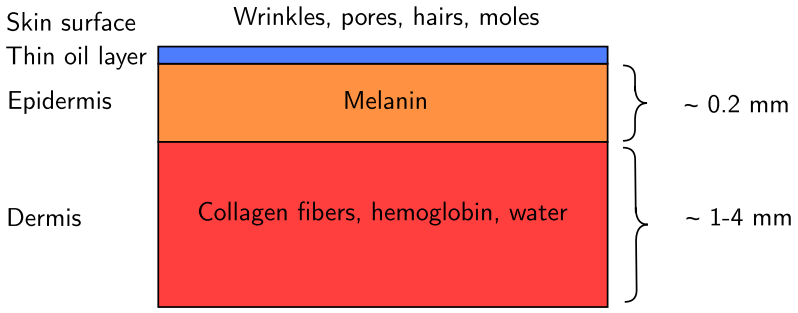

Figure 3.7: This figure shows a simplified model of the different layers of skin which is usually used in a model of skin for rendering.

3.3.3 Isotropic assumption

In our implementation of real-time subsurface scattering, we make the assumption that skin is isotropic. This means that for a given ray of light, the incoming azimuthal direction, relative to the surface normal, does not impact the visual appearance. In addition to this, it is defined as a semi-infinite homogeneous medium i.e., light will transmit into a medium which is homogeneous and infinite in one direction. The consequence of this is that only a single diffusion profile will be utilized. This is indeed a simplified version of human skin but it is often used in the context of real-time rendering. However, Donner et al. [7] uses a three-layer model, i.e., they use a composition of three diffusion profiles which yields a more accurate result. In the case for this thesis, we have omitted to generate more than a single diffusion profile. This is because it is not trivial for us to separate the diffusion phenomenon for each layer in skin with the current measurement setup.

3.3.4 Near-infrared light

In the electromagnetic spectrum, the spectrum which we can interpret with our eyes is known as the visible spectrum. This electromagnetic energy exists in the

wavelength range between 400-700 nanometers (nm). As for near-infrared light, it is commonly defined in the wavelength range between 700-2500 nm, see Figure 3.8. This wavelength range cannot be interpreted by the human eye and is invisible to us in normal conditions.

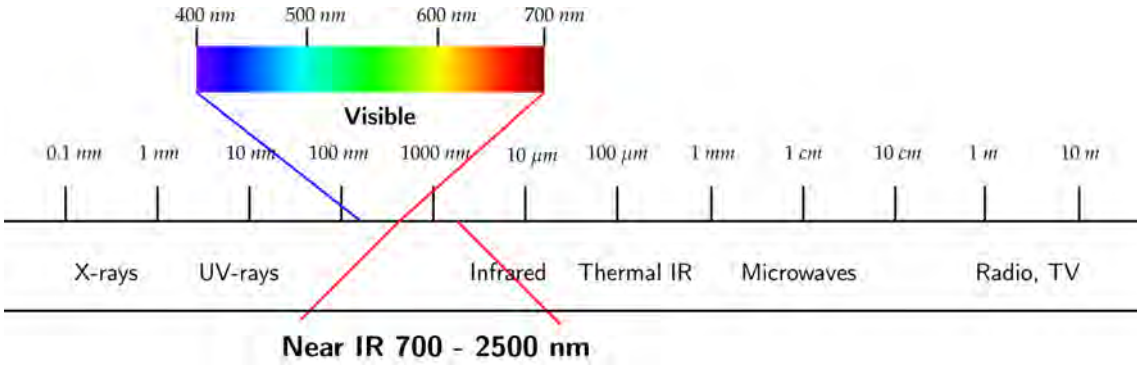

Figure 3.8: Basic representation of electromagnetic spectrum, defined on the wavelength axis.

Figure 3.9: Left images are images of a face illuminated by near-infrared (940 nm) light while the right column of images are illuminated in visible light and then grayscaled by the camera.

To produce this comparison we have used a single camera which can capture both IR and visible light. The camera settings are constant between the images. We utilize two light sources, a near-infrared light source and a LED light source. It should

be noted that we are not sure of the exact wavelength spectrum the light source in the visible spectrum covers but it should suffice to highlight the differences. In Figure 3.9 there is a notable difference between the visible and near-infrared images. In the near-infrared images, the skin appears more uniform. Some details become less obvious while others become more apparent. For instance, the blood vessels around the eye and the depths of the hair strands become more visible in near-infrared. This is mainly due to the reduced absorption of various features in the face. This is highlighted in the moles on the left cheek which do not seem to absorb near-infrared light, thus causing it to blend together with the skin. Further, it is easy to note other features behaving in a similar manner, e.g., the stubble, eyebrows and in general the skin appears brighter, implying a higher reflectance. The main reason for this behavior is because the face is illuminated by a singular wavelength. This wavelength shows less absorption in the details of the face which are normally highlighted by visible light through absorption. This causes the facial features to instead blend together in the skin, making it appear more uniform. Finally, the subsurface scattering effect appears more pronounced, in the sense that the skin appears smoother in areas such as the upper-lip, chin and cheeks.

3.4 Screen-space subsurface scattering

The are two state of the art methods for fast real-time subsurface scattering. The first method is the texture-space diffusion approach by D'Eon et al. [4]. The other approach, which is in nature very similar, is the screen-space diffusion approach by Jimenez et al. [12]. Our implementation relies on the latter and there are a number of reasons for this. First, the texture-space method does not scale well to multiple targets in the scene. For instance, if the scene contains multiple objects which require subsurface scattering, then each object must be processed separately in texture-space. On the contrary, the screen-space approach is indifferent to the number of targets. This is because the screen-space method is applied to a single image. By instead using a stencil buffer, or a mask of the scene with the rendering targets, it is possible to process the desired pixels. Furthermore, the texture-space diffusion method will process the target texture regardless of the distance from the camera. This is indeed inefficient as the detail of objects far from the camera are less noticeable. Again, this is not an issue in the screen-space approach as the depth buffer is used to test against a given threshold whether a pixel should be processed or not. A downside of using the screen-space approach is the difficulty of implementing translucency, i.e., light traversing *through* an object. This is due to the lack of information of back-facing geometry in screen-space. On the other hand, the texture-space diffusion method has no issues with this which is one of its upsides. In order to implement this phenomenon in the screen-space approach it would require additional techniques. However, under normal lighting conditions, the visual impact is almost negligible, thus it was left for future work. It was also deemed more trivial to implement the screen-space method in the existing pipeline. In this case it implemented as a single post-processing filter. Finally, Jimenez et al. conducted a small study to compare the differences between the two algorithms. They conclude that the visual differences between the texture-space

and screen-space are very minor and almost negligible [11]. In short, this method for approximating subsurface scattering, will operate on a rendered image and will effectively be a post-processing effect on the image. In screen-space the information is limited which means that we must store the necessary data and make it available in the final shader.

3.4.1 Kernel convolution

Because the chosen method operates in screen-space, we will be processing images, or textures. Ultimately, our main operation will be texture convolution with Gaussian filters of varying width. The act of convolution on an image is shown in Figure 3.10. In this case, a component-wise multiplication with the filter is shown. The respective products are summed and then becomes a new pixel in the output image.

The goal in the implementation of screen-space subsurface scattering is to create a blurring, or smoothing, effect of the target surface. The common filter for this purpose is a Gaussian filter, which is a common averaging filter. The primary reason for a Gaussian filter is the fact that it is radially symmetric. This means it can be used as a separable convolution. This follows from the the associative property of convolution:

$$f * (h * v) = (f * h) * v \tag{3.6}$$

where $*$ is the convolution operator, f is some input image and h, v are kernels. Thus, it can first be used as a horizontal filter, then subsequently as a vertical filter. This is an efficient method on modern GPUs but also because the computational complexity is reduced from $O(n^2)$ to $O(n)$. This is because a 2D filter requires a double loop to iterate and perform multiplications with the kernel and corresponding pixel [9].

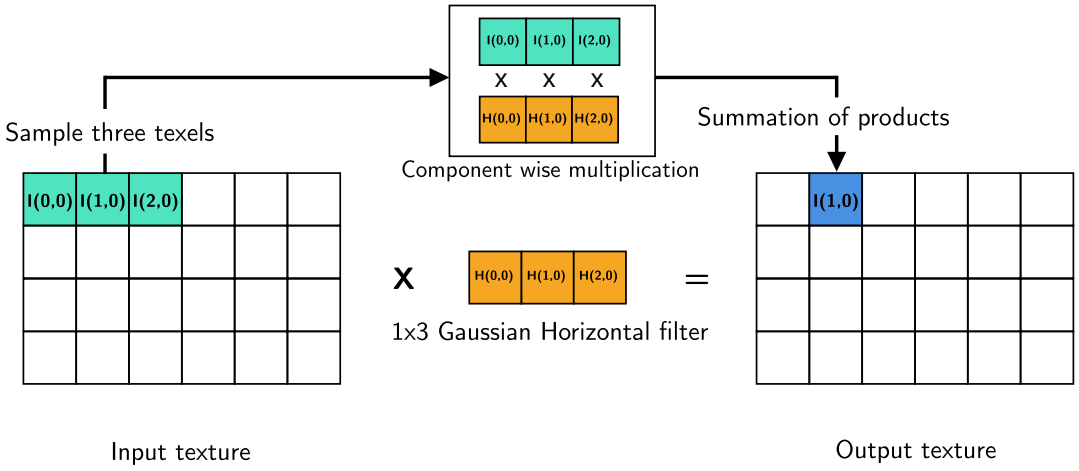

Figure 3.10: Texture convolution. A 1x3 filter is multiplied component-wise by a 1x3 sample of the the texture, then summed together into a new texture which does not change the size of the original texture.

The real-time method of implementing subsurface scattering is in short: smoothing the high-frequency content of the facial features while adjusting the kernel width according to the curvature of the surface. As already mentioned we are going to use the common Gaussian filter for this purpose. In order to generate a Gaussian filter, a Gaussian function is sampled at discrete intervals, e.g., $-2, -1, 0, 1, 2$, which would be a 5-sample wide kernel. If we observe a typical zero-mean Gaussian function, the largest weights will appear in the center of the function. As we deviate further from the mean of the Gaussian, the smaller the values and consequently, these samples will carry less weight in the final summation. Each sample from a Gaussian will function as a weight to be multiplied with a pixel. This process can be described by the following equation:

$$g(x) = \omega * f(x) = \sum_{dx=-a}^{a} \omega(dx) f(x + dx) \tag{3.7}$$

where ω is the Gaussian kernel in this case, dx are the discrete interval steps, $f(x)$ is the current pixel and$-a$ to a is the discrete sampling interval.

3.4.2 Adjusting kernel width

When applying the smoothing filters to the target surfaces, it is important that the amount of smoothing is adjusted correctly. We must adjust the impact of the filtering process depending on the surface geometry. For instance, surfaces which have a large angle between the surface normal and the camera position should be less impact on the amount of smoothing. This can be achieved with multiple methods. It is possible to use normal map of the object to calculate when the angle between camera and surface is large. Another approach is to use the depth buffer to inspect when there is a large depth difference between two pixels. The latter method was chosen for this thesis. Ultimately, the computed depth gradient value can be used to adjust the kernel width, or height, of the Gaussian filter.

When applying this method to the final image in post-processing, the kernel size must be adjusted in an adaptive manner. The major reason for this is so that the low-frequency content is not overly blurred. The smoothing should be applied in areas which are relatively smooth, e.g., the cheeks and the forehead, and be less aggressive in other areas such as the nose. Jimenez et al. [11] compute stretch factors which uses a derivative kernel to compute the depth difference between two pixels. The equations for the horizontal and vertical stretch factor are computed in the following manner:

$$s_x = \frac{ssslevel}{d(x,y) + correction \cdot min(abs(ddx(d(x,y), maxdd)))} \tag{3.8}$$

$$s_y = \frac{ssslevel}{d(x,y) + correction \cdot min(abs(ddy(d(x,y), maxdd)))} \tag{3.9}$$

$d(x, y)$ is the depth of the pixel at (x, y), given by a linear depth map, ddx and ddy are built in functions in the *High-level Shading Language* (HLSL) [5], which returns

the depth gradient. "ssslevel", "correction", and "maxdd" are constant parameters which are found empirically for a given scene. They control how aggressive the algorithm should be when smoothing low-frequency content, i.e., how much of the depth gradient should be taken into consideration as a scaling factor for the kernel width. This stretch factor is then multiplied with the original size of the kernel as a way of scaling the kernel.

In Figure 3.11 the effect of the parameters *correction* and *ssslevel* are shown explicitly, allowing for control of the subsurface approximation effect. Note that some of the parameter values are excessive but demonstrates their effect.

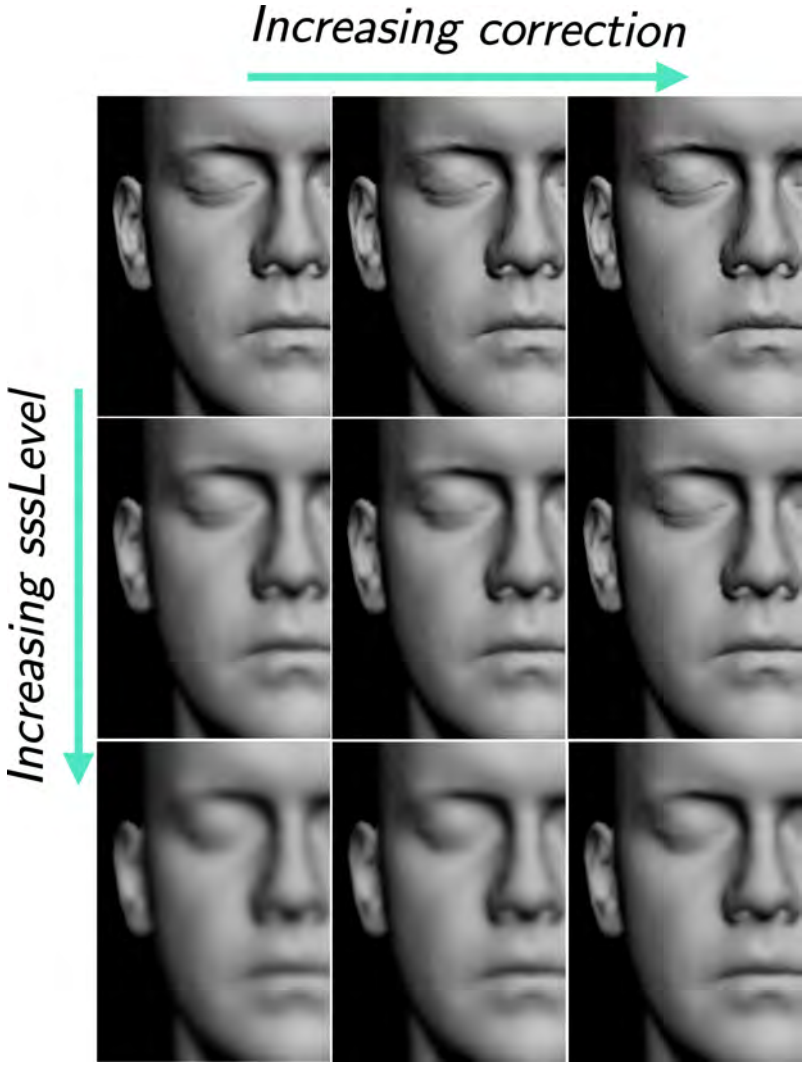

Figure 3.11: This matrix of images demonstrates the influence of the *ssslevel* and *correction* parameters. From top to bottom $ssslevel = 0.3, 0.6, 1.0$ and from left to right $correction = 800, 1500, 2500$. Note how as the correction parameter is increased, the subsurface scattering effect decreases in low-frequency areas, by increasing ssslevel, the subsurface scattering effect increases.

4

Method and Implementation

This chapter demonstrates our method to measure the subsurface scattering profile of near-infrared light with skin. In addition, an implementation of real-time subsurface scattering is shown. By using the data acquired from the experiment, a diffusion profile of near-infrared light and human skin is created by fitting a sum-of-Gaussians to the data. This profile is used in an implementation of screen-space subsurface scattering.

4.1 Measure of subsurface scattering

This experiment is inspired by the measuring method presented by Jensen et al. [10]. Their method of experiment captures the scattering properties of a translucent material where the scattering coefficient dominates.

4.1.1 Measurement setup

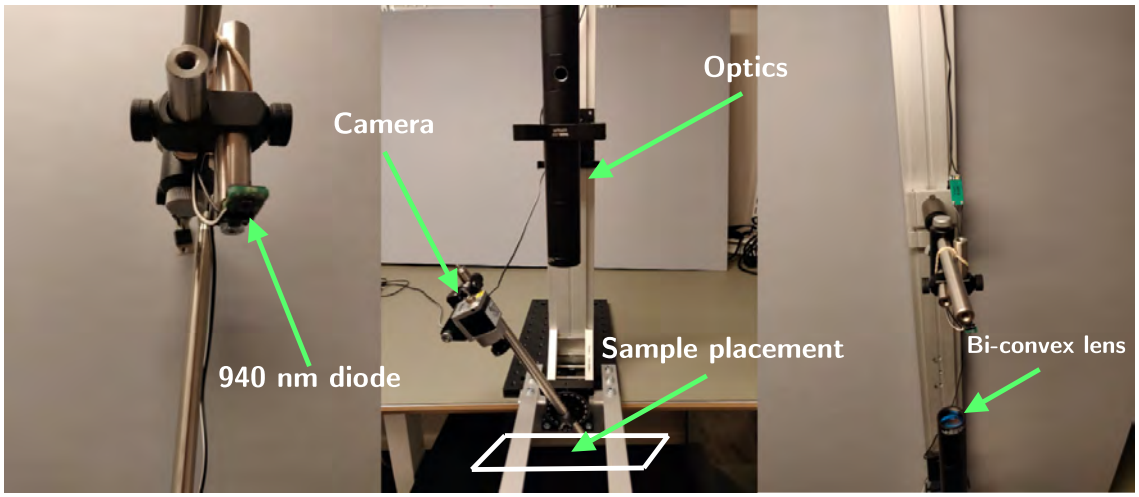

Figure 4.1: Measuring setup. Three images of the same measuring device. We measure the diffusion profile by using a set of optics, a near-infrared 940 nm diode and a camera which can capture light outside the visible spectrum.

In the experiment, an image is captured as the skin is illuminated by the near-infrared diode. This means that the evaluation is based on the pixel intensity output produced by the camera. In order for the pixel values in the image to be meaningful,

it is required that the measurements have a reference to be compared against. If no such reference exists, the pixel values in the image become arbitrary since the pixel intensity is highly dependent on the settings of the camera. When performing the experiment a calibrated material was used. A calibrated material is one where the amount of light reflecting back from the surface is known. A common such material, which is also used in the measurement method by Jensen et al. is Spectralon. Note that the calibration of the material is not conducted by us. We are able to use Spectralon as a reference because human skin is a highly scattering material. When light has entered into the skin and scattered a number of times, one can assume its exiting direction will be entirely random. As the ray exits its behavior is similar to the phenomenon of diffuse scattering. This assumption allows for all the measurements to be relative to the Spectralon surface as it is a calibrated diffuse surface. In our case, our Spectralon surface is measured to reflect more than 99% of all incoming light diffusely. In this order, we can deduce the reflectance of a given pixel in an image by capturing two images, one of the Spectralon surface, and another of human skin.

Figure 4.2: Spectralon material, calibrated to reflect more than 0.99 % of incoming light diffusely.

In this experiment, the objective was to record the exponential falloff curve of energy as it scatters on the surface of the skin. We focus a light source into a small point on the surface of the skin and measure the intensity falloff by observing pixel values in the resulting image.

For the experiment, the measuring setup used is shown in Figure 4.1. It uses a single light source diode which emits near-infrared light at 940 nm. Further, a Basler aCA1920-155μ camera with a 25 mm lens was used. The camera also uses a bandpass filter which limits the incoming wavelength spectrum to 940-950 nm. Additionally, a configuration of lenses was used to focus a partial part of the near-infrared light into a smaller point. This optical setup consisted of two bi-convex lenses on both ends of an enclosing cylinder. This effectively resulted in a size-reduction factor of two of the actual diode light chip. The final focused point of near-infrared light is approximately 1 mm^2. The camera is positioned on a rotatable axis. This way, the angle relative to the surface on the skin could be adjusted

to avoid specular contributions from the top layer of the skin.

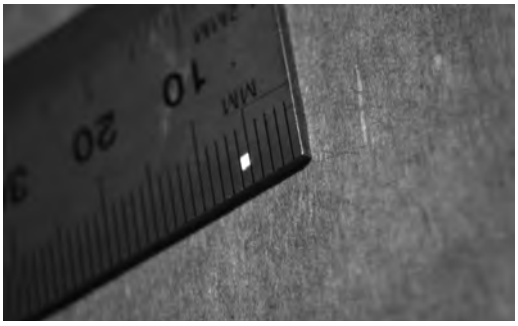

Figure 4.3: Measured size of the diode image in the plane of focus. The diode light image becomes approximately 1 mm in both horizontal and vertical direction.

4.1.2 Measurement of subsurface scattering

The experiment began by capturing an image of the Spectralon surface with an exposure time such that there are **no** overexposed pixels. If there are overexposed pixels then we would lose our reference of energy, and again the pixel values become arbitrary. The exposure time of the camera is adjusted such that the average pixel intensity of illuminated area was around 240-250. Note that the camera used 8 bits to store each pixel value. This means that the intensity values range between 0-255.

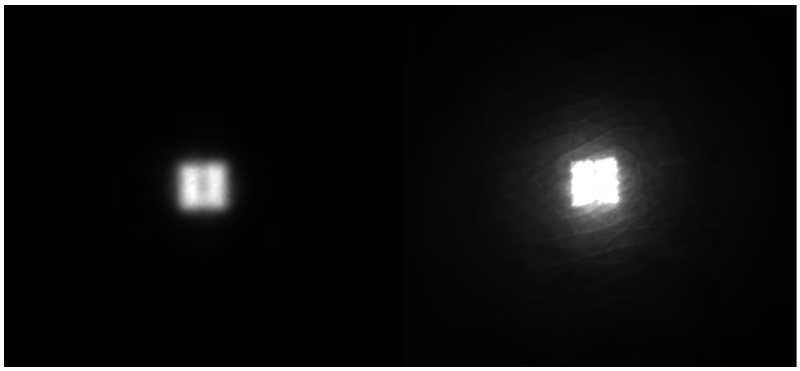

Figure 4.4: Left: Spectralon illuminated by near-infrared light. **Right**: Inner forearm skin illuminated by near-infrared light. Both are using the same exposure time and we can see noticeable subsurface scattering in the right image.

In Figure 4.4, Spectralon and skin is illuminated by near-infrared light. In the image with skin there are many overexposed pixels, however, the overexposed pixels are not used to measure the subsurface scattering profile. In this case they do not matter because the goal is to separate the multiple subsurface scattering effect from any specular or single-scattering effects. Thus, the only pixels which are relevant are the pixels which exists outside the square area of direct illumination. The area of the image which is used to generate the diffusion profile is shown in Figure 4.5. In this image, the measurements start from the center of the illuminated area as is

done by Jensen et al. [10]. The diffusion profile is generated from this image by averaging each column of the sample area. This creates a 1D profile of the energy falloff representing the subsurface scattering profile, the diffusion profile.

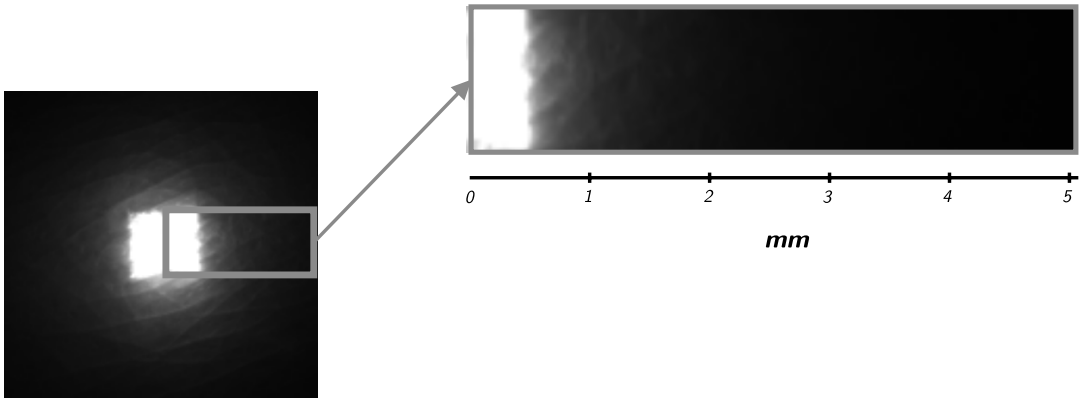

Figure 4.5: Subsurface scattering for skin illuminated by the light source at 940-950 nm. A 5 mm wide rectangular sample is taken and for each column we generate an average pixel intensity.

After gathering the data, it must be adjusted with respect to the average pixel intensity of the Spectralon surface. Additionally, in this experiment, an optical filter producing a light reduction of a factor 10 was used which must be taken into account. In the experiment, the average pixel intensity for the Spectralon surface illuminated by near-infrared light was approximately 244, combined with the reduction factor of 10. Thus, each data point in the 1D diffusion profile of skin is multiplied by 1/2440. The result from this adjustment can be seen in Figure 4.6.

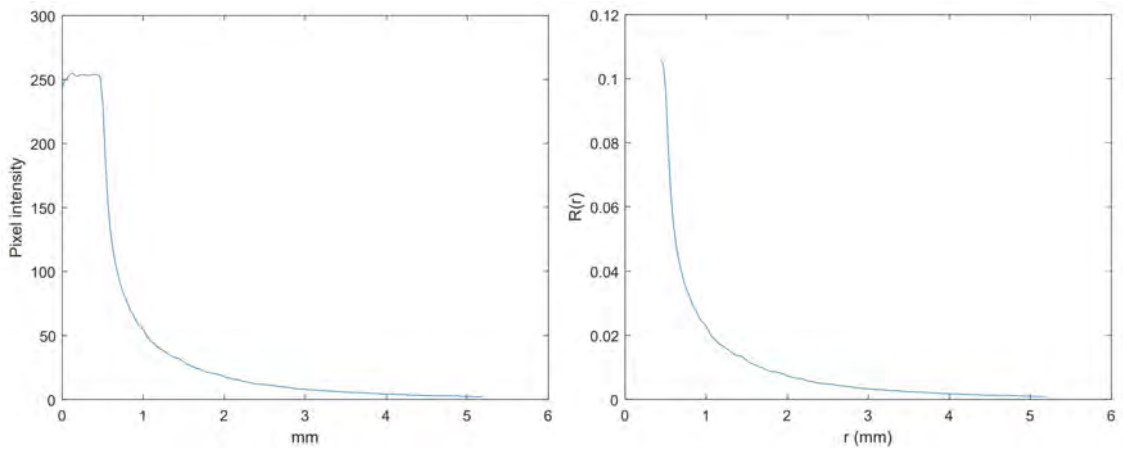

Figure 4.6: Left: Raw data from the sample area in Figure 4.5. **Right**: Data adjusted to the average Spectralon pixel intensity and the reduction factor from the optical filter. The Y-axis now represents the diffuse surface reflectance, **R(r)**.

4.1.3 Function fitting

To fit a function to the data we used a sum-of-Gaussians per D'Eon et al. [4]. To approximate the data, we utilized existing curve fitting tools in Matlab [16], with the following definition for the Gaussian function:

$$G(v, r) = w \cdot \frac{1}{2\pi v} e^{(-r^2/(2v))} \tag{4.1}$$

where v is the variance, r is the distance, and w is the weight. The final expression for the diffusion profile is given by

$$\sum_{i}^{k} w_i G(v_i, r) \tag{4.2}$$

The final fitted function is presented in the results section with corresponding weights and variance parameters.

4.2 Screen-space subsurface scattering

The chosen algorithm for real-time subsurface scattering in this thesis is the screen-space approach as presented by Jimenez et al. [11]. This algorithm can be summarized as a sequential number of shader passes of a rendered image. This shader typically operates in a post-processing stage where the smoothing of the desired objects are performed with, in our case, five unique Gaussian kernels. As mentioned in the theory section, this shader should not smooth uniformly across the mesh but must also take into account the geometry of the object to decides the degree of smoothing.

Implementing the width-adjusting kernel is performed by using a derivative kernel to differentiate between depth differences on the mesh. Coupling this with a reasonable chosen threshold, the size of the Gaussian blur filter can be adjusted across the mesh. Finding the derivative of the depth buffer is made easy by using built-in functions in the shader language, *ddx* and *ddy*. By sampling the depth buffer, these functions return the depth difference between two pixels in the horizontal or vertical direction. The effect of this filter applied to the input texture is shown in Figure 4.7 with a threshold value quite low, producing an exaggerated effect.

(a) Face model.

(b) Face model with derivative kernel.

Figure 4.7: Left: Rendered face model. **Right**: White area represents the surfaces where the gradient is relatively large thus allowing us to adjust the Gaussian kernel size accordingly in those areas.

4.2.1 Mask texture

In the original paper, a stencil buffer is used to separate the surfaces which are supposed to have subsurface scattering. Due to some difficulty of using the stencil buffer in our application in the Unity engine we opted to simply add an additional render pass. In this pass, the meshes which are to include subsurface scattering are rendered into the mask texture and subsequently the texture is saved as a global texture. In order for the post-processing shader to not affect any occluding objects, the depth buffer of the camera and the local depth (after transforming its depth values into camera space) of the object are compared to deduce what geometry is occluding the target mesh, see Figure 4.8.

(a) Face model and occluding sphere **(b)** Face model mask with occluding object taken into consideration

Figure 4.8: Face model is occluded by a sphere..

4.2.2 Post-processing shader

The listing below is the shader written using Unity's shader language. The shader itself is very similar to a standard separable Gaussian blur shader. The exception is that we compute a stretch factor which adjusts the kernel width and performs a check against the mesh mask to determine if the post processing filter should be applied to the given texels.

```
1
2 TEXTURE2D_SAMPLER2D(_MainTex, sampler_MainTex);
3 TEXTURE2D_SAMPLER2D(_CameraDepthTexture, sampler_CameraDepthTexture
    );
4 TEXTURE2D_SAMPLER2D(_MeshMask, sampler_MeshMask);
5
6 float sssLevel;
7 float correction;
8 float kernel_width = 7.0f / _ScreenParams.x;
9
10
11 float4 FragH1(VaryingsDefault i) : SV_Target
12 {
13     float4 t = SAMPLE_TEXTURE2D(_MeshMask, sampler_MeshMask, i.
    texcoord);
14
15     float w[7] = {0.1053, 0.1402, 0.1664, 0.1762, 0.1664, 0.1402,
    0.1053};
16
```

```
17      float depth = LinearEyeDepth(SAMPLE_DEPTH_TEXTURE(
        _CameraDepthTexture, sampler_CameraDepthTexture, i.
        texcoordStereo));
18
19      if(t.x == 0.0 && t.y == 0.0 && t.z == 0.0)
20          return SAMPLE_TEXTURE2D(_MainTex, sampler_MainTex, i.
        texcoord);
21
22      // Compute stretch factor
23      float s_x = sssLevel / (depth + correction * min(abs(ddx(depth)
        ), maxdd);
24
25      // Multiply stretch factor
26      float2 finalWidth = s_x * kernel_width * float2(1.0, 0.0);
27
28      // Subtract offset from current texture coordinates.
29      float2 offset = i.texcoord - finalWidth
30
31      float4 color = float4(0.0, 0.0, 0.0, 1.0);
32
33      for(int i = 0; i < 7; i++){
34          float3 tap = SAMPLE_TEXTURE2D(_MainTex, sampler_MainTex,
        offset);
35          color.rgb += tap * w[i];
36          offset += finalWidth / 3.0;
37      }
38
39      return color;
40
41  }
```

Listing 4.1: Horizontal Gaussian shader written in Unity's shader language.

In Listing 4.1, *MainTex* is the rendered image that will be blurred by the Gaussian, *CameraDepthTexture* is the depth buffer from the camera, *MeshMask* is the mask texture of the target mesh. This shader is one of the horizontal Gaussian passes which occur, and the vertical shader is almost identical.

The parameters *sssLevel* and *correction* are parameters to control the level of smoothing. Both of the parameters are set empirically and have no real physical meaning. The *correction* parameter controls the influence of the depth derivative kernel, e.g., when the difference between two pixels is large and a large value is multiplied with the *correction* value, this effectively divides the *sssLevel* value by a relatively large number. This effectively reduces the impact from the stretch factor. Conversely, if the depth difference between two pixels is very small, then the correction parameter has essentially no impact at all. Thus, the goal is to have the depth impact be minimal at smooth surfaces and high impact in areas where many changes occur, e.g., around the nose, edges of the head and mouth.

At **line 23**, the stretching factor is computed and then multiplied at **line 26** with the original kernel width and a horizontal axis vector. This way, only the horizontal axis is affected.

In **lines 33-37**, the sampling of the rendered texture is performed. The division by 3.0 is because we offset the current texture coordinates by *finalWidth*. In total, seven samples are taken, three are to the left of the original texture coordinates. Then we sample the original texture coordinates, finally we sample three to the right. This is further clarified in Figure 4.9. At each sample, the sample is multiplied by its corresponding weight defined at **line 15**.

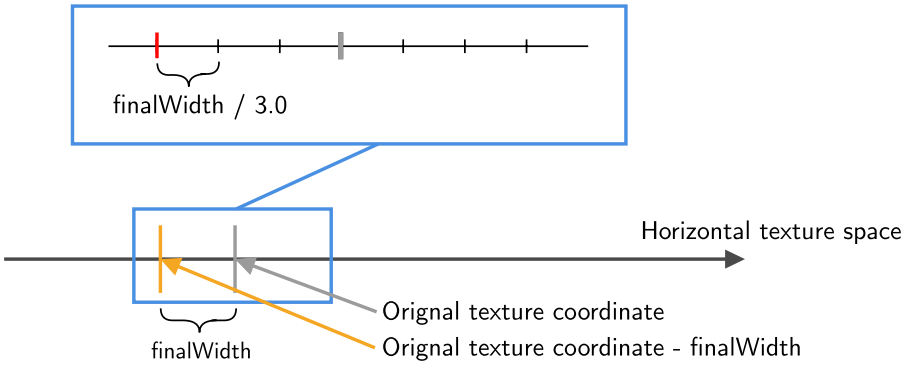

Figure 4.9: *finalWidth* is subtracted from the original texture coordinates, then the samples move step by step in intervals of *finalWidth* / 3.0 to produce a total of 6 samples around the original texture coordinates.

The final shader, which is a sum of five Gaussians, requires five subsequent horizontal passes followed by five vertical passes. The input texture is convolved with a Gaussian kernel, then in the following pass, convolved again, but with a different Gaussian and so on, see Figure 4.10. Five unique Gaussians are applied for both the horizontal and vertical direction. The end result is a mesh which is filtered in *ten* passes, ultimately blurring the desired surfaces of the target mesh.

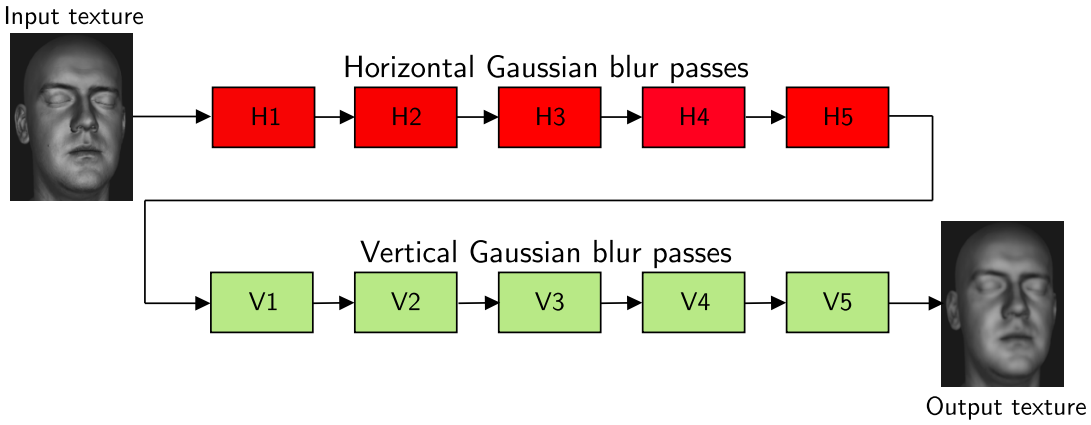

Figure 4.10: A given input image is passed onto several horizontal and vertical Gaussian blur passes in a sequential manner.

4.3 Finding correct parameters for near-infrared light

There is some ambiguity in how to determine the parameters which are used in the algorithm. Even though we have many filters applied to the image that generate a smoothing effect, the parameters determining the stretching factors also play a large role in the degree of smoothing applied. Initially, an attempt was made to measure these parameters by conducting a similar experiment to that in Section 4.1 inside the engine. However, we were not able to conduct this experiment in a physically correct manner at the time. In order to make an accurate experiment we would need to correctly control the exposure of the in-application camera, the tonemapping, the light intensity and so on. Instead we opted to determine these parameters empirically by comparing our results to real-life images and adjust the parameters such that similar behavior is achieved.

5
Results

In this section we demonstrate the results. This includes our data acquired from the measurements and the corresponding sum-of-Gaussian fit. Additionally we compare this to diffusion profiles of skin produced by Jensen et al. [10]. Finally we show in-engine results of the shader applied to a face model with parameters gathered from Section 4.3.

5.1 Diffusion profile

In Figure 5.1, we show the final sum-of-Gaussian approximation made from five zero-mean Gaussian functions with the parameters listed in Table 5.1. The invididual Gaussians functions are shown in Figure 5.2.

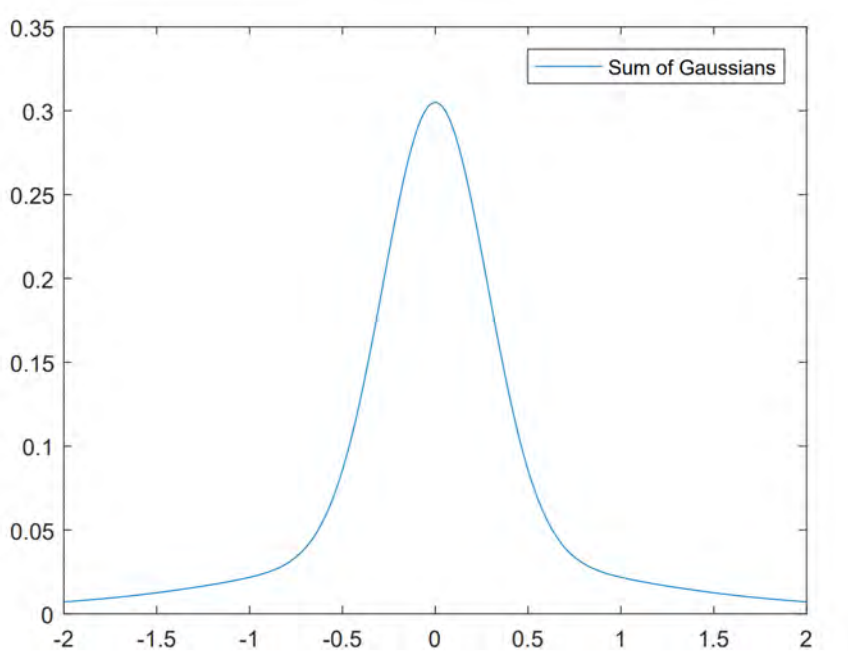

Figure 5.1: Sum of six zero-mean Gaussian functions.

Table 5.1: Weights and variance parameters

Gaussian function parameters and weights					
v_i	8,749	2.099	0.549	0.5574	1.194
w_i	0.1478	0.1323	0.0004	0.0527	0.090

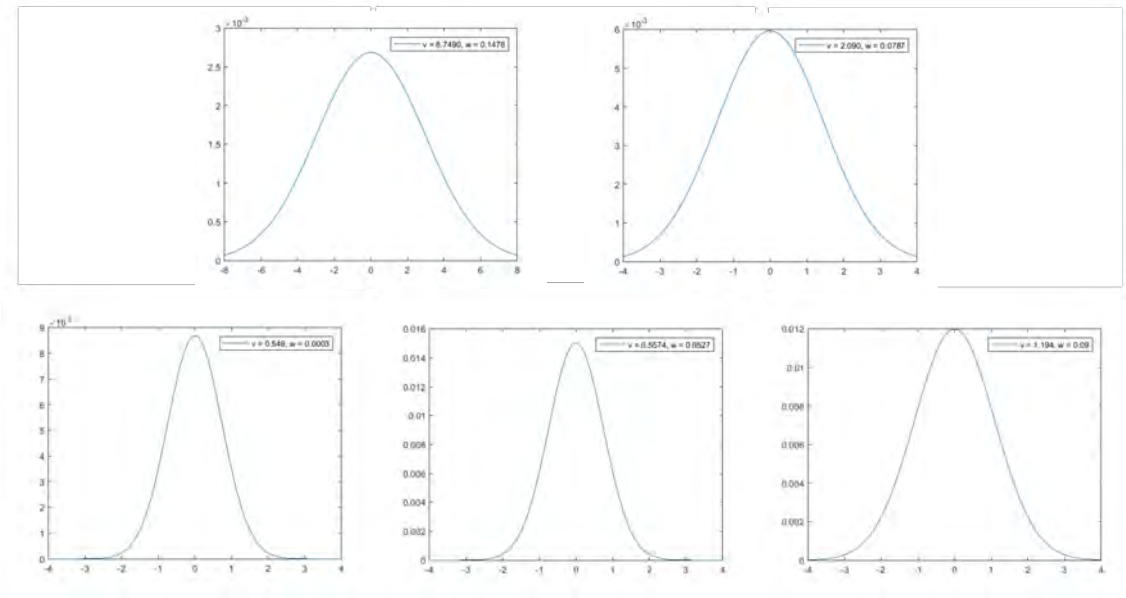

Figure 5.2: Individual Gaussian functions.

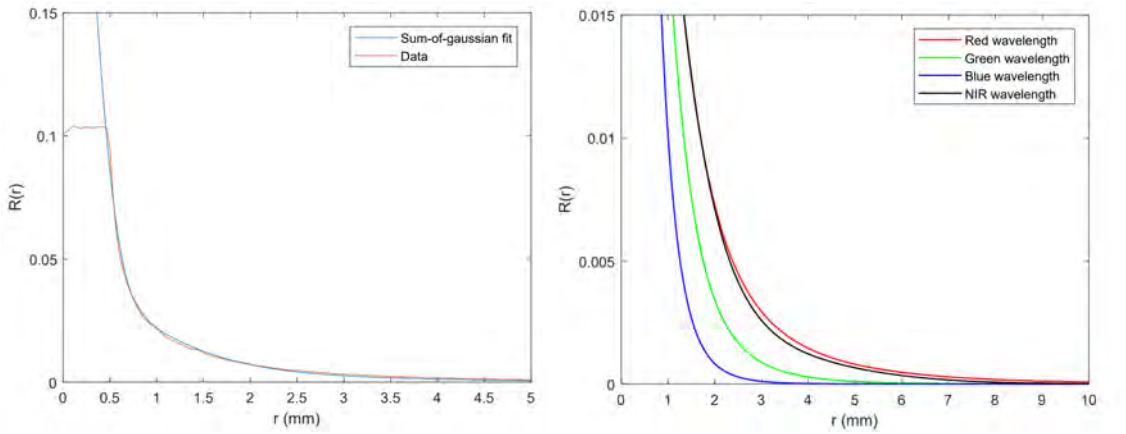

Figure 5.3: **Left**: Sum-of-Gaussian approximation against the data gathered from the experiment. **Right**: Sum-of-Gaussian function fit (**black line**) against red, green and blue wavelengths. All profiles are measured against human skin.

5.2 Real-time approximation

5.2.1 Subsurface scattering approximation

The following image demonstrates the combination of diffuse, specular and the implemented screen-space subsurface scattering into a single image. Note that for the diffuse and specular lighting, the standard shader in Unity is utilized. The subsurface scattering is directly applied to the result of the diffuse and specular component.

(a) Diffuse **(b)** Specular **(c)** Subsurface scattering **(d)** Final combined result

Figure 5.4: This figure shows the separate components and how they are added together for a final result. In this example we use following parameters, $ssslevel = 0.5$, $correction = 2000$.

5.2.2 Comparing near-infrared scattering with visible wavelengths

In Figure 5.5, the parameters *ssslevel* and *correction* are kept constant in order to demonstrate the difference in smoothing between the different wavelengths. Additionally, we only render half of the image with our smoothing filter in order to highlight the difference between the non-smoothed surface and the smoothed surface. We compare the level of smoothing applied by the visible wavelengths, red, green and blue against near-infrared smoothing. We can observe that red and near-infrared behave similarly but green and blue do show significantly less smoothing which is accurate to the diffusion profiles shown in Figure 5.3. The Gaussian parameters for red, green and blue wavelengths are gathered from D'Eon et al. [4].

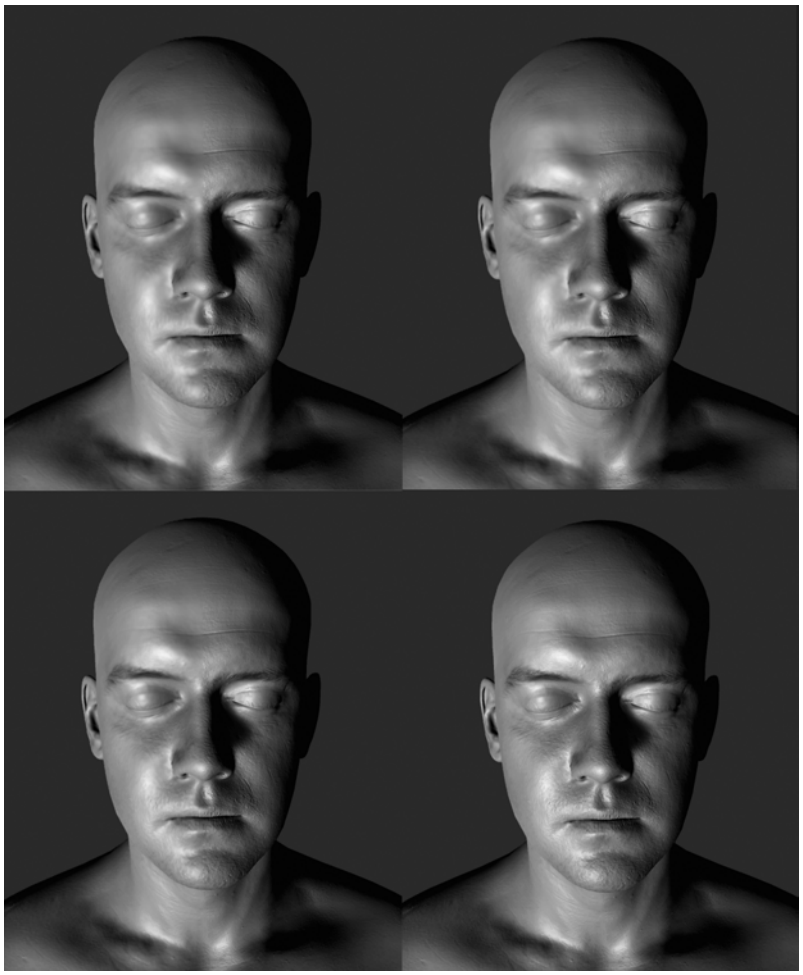

Figure 5.5: Top left: Near-infrared diffusion profile. **Top right**: Red diffusion profile. **Bottom left**: Green diffusion profile. **Bottom right**: Blue diffusion profile. Parameters are constant in all images, $ssslevel = 0.4$, $correction = 4500$.

In this final figure we show the face model without any color grading and smoothing and then compare this against our final result. The parameters are chosen empirically, by looking at real life images of faces illuminated by near-infrared light. We also attempt to adjust the color grading of the image such that it behaves approximately like near-infrared light, however, this is also simply based on observation of real-life images and it is obviously not entirely accurate.

Figure 5.6: Left: No color grading or subsurface scattering. **Middle**: Subsurface scattering applied. **Right**: Subsurface scattering and color grading. $ssslevel = 0.7$, $correction = 7500$.

6
Discussion and Conclusions

6.1 Discussion

This section discusses the results of the thesis. We discuss the measurements, near-infrared light and skin and the implementation of screen-space subsurface scattering. Finally we present our conclusions of the thesis.

6.1.1 Measurements

The measurements produced by the experiment are reasonable in the sense that the results behaved similar to that of red light, which is expected. However, they are inconclusive because our data is directly compared to the measurements made by Jensen et al. [10]. These measurements may not have been conducted under the same exact circumstances as our own. For instance, in this thesis we measured Caucasian skin on the inside of the forearm and it is likely that we did not measure at the same location or on the same type skin as they did. Further, it would have been beneficial if the measurements of red light and near-infrared were conducted simultaneously such that they could be compared directly. If given future opportunity to redo the experiment, it will be conducted on multiple people, at multiple positions on the body and further it would have been done with both red light and near-infrared light such that it would be possible to compare against our own measurements.

6.1.2 Near-infrared light and human skin

The comparison between the person in Figure 3.9 show very significant differences which are not specifically tackled in the thesis. In the comparisons, the Lee-Perry Smith 3D head scan [17] was used. This model uses an albedo texture which is generated by scan of a person illuminated by light in the visible spectrum. This face scan is then transformed into an RGB albedo texture. The realization, in hindsight, is that the uniform appearance of skin in near-infrared light is also due to the restricted spectrum of wavelengths. The visible wavelengths which are otherwise absorbed in the different parts of the skin, giving the various details of skin their color, are not present in near-infrared wavelengths which results in a more uniform appearance. Ideally, an albedo texture, or rather, a 3D scan of a face should be constructed, where the face is illuminated by near-infrared light. It should also be mentioned that the target application where the smoothing was implemented, had a number of different types of face models, each of which could be manipulated freely.

This allows the developer to test how a tracking algorithm behaves with a certain type of face. This also means that simply generating a 3D scan of a face illuminated by near-infrared light might not be the solution for the target application, rather something more general would need to be applied.

As previously mentioned, although inconclusive, it is likely, that the subsurface scattering effect is increased slightly but is not the only factor to the increased smoothness of the skin which was the hypothesis when the thesis began.

6.2 Conclusion

The original thesis question was to investigate if the degree subsurface scattering of near-infrared light in human skin was greater than that of light in the visible spectrum. The conclusion of the results is that near-infrared light scatters more than blue and green wavelengths. Further, it scatters approximately as much as red light but we are unable to make any conclusions if it scatters more. To make a conclusive statement it requires us to make more measurements. These measurements would be at different locations but also for different skin types. Additionally, similar experiments should be conducted for the typical red, green and blue wavelengths in order to more concretely gauge the accuracy of our measurement method.

It should also be noted that the measurement technique adopted from [10], which is originally a method applied in the visible spectrum, could also sucessfully be applied outside the visible spectrum.

6.2.1 Ethical considerations

We do not see any direct implications from the work conducted in this thesis. But we can conclude that a basic purpose for understanding material properties is to invent techniques that can be used to reduce the gap between reality and simulation. This can of course be used in a wide number of areas and industries. An example is that with sufficiently real environments, experiments which require visual accuracy can be performed with high accuracy and subsequently be directly transferred to reality without endangering humans. Thus, the work in this thesis will, at most, only have a minor indirect implication by its existence.

6.2.2 Future work

In regards to measurements, it would be interesting to produce more data for Caucasian skin but also for other types. We would also like to conduct measurements with red light and human skin and compare this with our own measurements. Additionally it would be interesting, to further verify the correctness, to measure other types of materials, e.g., marble and make sure that our method is still valid.

Another topic to explore are methods of transforming the RGB albedo textures into near-infrared. In order to generate accurate albedo textures, it is necessary to

either generate a 3D scan of a face where the face is illuminated by near-infrared light or attempt to approximate the effect of near-infrared illumination by other methods. Although, we have only briefly investigated this topic, Reiter et al. [23] present a method where they use an RGB input of a face and are able to transform it into a texture with the skin behaving like it is being illuminated by near-infrared light.

Bibliography

[1] R Rox Anderson, John A Parrish, et al. "The optics of human skin." In: *Journal of investigative dermatology* 77.1 (1981), pp. 13–19.

[2] Gyeongmin Choe, Srinivasa G Narasimhan, and In So Kweon. "Simultaneous estimation of near ir brdf and fine-scale surface geometry". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016, pp. 2452–2460.

[3] Catherine C Cooksey, Benjamin K Tsai, and David W Allen. "Spectral reflectance variability of skin and attributing factors". In: *Radar Sensor Technology XIX; and Active and Passive Signatures VI*. Vol. 9461. International Society for Optics and Photonics. 2015, p. 94611M.

[4] Eugene d'Eon and David Luebke. "Advanced techniques for realistic real-time skin rendering". In: *GPU Gems* 3.3 (2007), pp. 293–347.

[5] Microsoft Documentation. *ddx*. URL: `https://docs.microsoft.com/en-us/windows/win32/direct3dhlsl/dx-graphics-hlsl-ddx` (visited on 03/23/2020).

[6] Craig Donner and Henrik Wann Jensen. "A Spectral BSSRDF for Shading Human Skin." In: *Rendering techniques* 2006 (2006), pp. 409–418.

[7] Craig Donner and Henrik Wann Jensen. "Light diffusion in multi-layered translucent materials". In: *ACM Transactions on Graphics (ToG)* 24.3 (2005), pp. 1032–1039.

[8] Takanori Igarashi, Ko Nishino, Shree K Nayar, et al. "The appearance of human skin: A survey". In: *Foundations and Trends® in Computer Graphics and Vision* 3.1 (2007), pp. 1–95.

[9] Intel. *An investigation of fast real-time GPU-based image blur algorithms*. 2014. URL: `https://software.intel.com/content/www/us/en/develop/blogs/an-investigation-of-fast-real-time-gpu-based-image-blur-algorithms.html` (visited on 06/18/2020).

[10] Henrik Wann Jensen et al. "A practical model for subsurface light transport". In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. 2001, pp. 511–518.

[11] Jorge Jimenez and Diego Gutierrez. "GPU Pro: Advanced Rendering Techniques, chapter Screen-Space Subsurface Scattering". In: *AK Peters Ltd* (2010), pp. 335–351.

[12] Jorge Jimenez and Diego Gutierrez. “Screen-space subsurface scattering”. In: *GPU Pro: Advanced Rendering Techniques* (2010), pp. 335–351.

[13] Jorge Jimenez et al. “Separable subsurface scattering”. In: *Computer Graphics Forum*. Vol. 34. 6. Wiley Online Library. 2015, pp. 188–197.

[14] Yusuke Kanzawa, Yoshikatsu Kimura, and Takashi Naito. “Human skin detection by visible and near-infrared imaging”. In: *IAPR Conference on Machine Vision Applications*. Vol. 12. Citeseer. 2011, pp. 14–22.

[15] Stan Z Li et al. “Highly accurate and fast face recognition using near infrared images”. In: *International Conference on Biometrics*. Springer. 2006, pp. 151–158.

[16] MathWorks MatLab. *Curve Fitting*. 2019. URL: `https://se.mathworks.com/help/curvefit/curvefitting-app.html` (visited on 05/11/2020).

[17] Morgan McGuire. *Computer Graphics Archive*. `https://casual-effects.com/data`. July 2017. URL: `https://casual-effects.com/data`.

[18] Michael J Mendenhall, Abel S Nunez, and Richard K Martin. “Human skin detection in the visible and near infrared”. In: *Applied optics* 54.35 (2015), pp. 10559–10570.

[19] Mastura Mohamad et al. “Correlation between near infrared spectroscopy and electrical techniques in measuring skin moisture content”. In: *Journal of Physics: Conference Series* 546 (Nov. 2014), p. 012021. DOI: `10.1088/1742-6596/546/1/012021`.

[20] Yasuhiro Mukaigawa, Kazuya Suzuki, and Yasushi Yagi. “Analysis of subsurface scattering based on dipole approximation”. In: *IPSJ Transactions on Computer Vision and Applications* 1 (2009), pp. 128–138.

[21] Eric Penner and George Borshukov. “Pre-integrated skin shading”. In: *Gpu Pro* 2 (2011), pp. 41–55.

[22] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016.

[23] Michael Reiter et al. “3D and infrared face reconstruction from RGB data using canonical correlation analysis”. In: *18th International Conference on Pattern Recognition (ICPR’06)*. Vol. 1. IEEE. 2006, pp. 425–428.

[24] Sarah Tariq et al. “Efficient estimation of spatially varying subsurface scattering parameters”. In: *Vision, Modeling, and Visualization (VMV2006)* (2006), pp. 129–136.

[25] Tim Weyrich et al. “Analysis of human faces using a measurement-based skin reflectance model”. In: *ACM Transactions on Graphics (TOG)* 25.3 (2006), pp. 1013–1024.

A

Appendix 1