AN ARTISTIC APPROACH FOR INTUITIVE CONTROL OF LIGHT TRANSFER IN PARTICIPATING MEDIA

A Thesis

by

JOSE BENJAMIN GUINEA MONTALVO

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

Approved by:

Chair of Committee,	Ergun Akleman
Committee Members,	Jon Keyser
	Ann McNamara
	Richard Davison
Department Head,	Tim McLauglin

December 2012

Major Subject: Visualization

ABSTRACT

The sole purpose of every form of visual representation is to make something look believable. Even among abstract or conceptual representation, the purpose is to create something that within the defined visual language the audience will consider believable and accepted.

In the field of computer generated representation there are numerous visual languages that have been developed throughout the years, attempting to solve different visualization or artistic problems. This thesis presents an alternative light transfer model for participating media focused on the intuitive control of the illumination data and the artistic value of the resulting image. The purpose is not focused on accurately modeling lights physical behavior and its interaction with the surfaces and elements.

My thesis describes an artistic approach which aims to offer an organic and intuitive control of the glow and temperature of the effects of participating media and direct the value and hues through the surfaces. The system described in the thesis approximates light transfer through a given volume by calculating light contribution in the volume with discreet sampling and subsequently gathering these values to determine the diffuse scattering contribution for the volume.

I will also discuss the assumptions made to allow such approximations, as well as how the intuitive control offered by the approach and these approximations allow new forms or representation and artistic direction.

DEDICATION

To my parents, Jesus and Lorena, for all the support and interest that they have always shown me. They are the root of my curiosity, interest and passion for challenges and solutions.

To my brothers, Sam, Pablo, Pietro and Gabriele. We span a couple of generations, but we share that interest and curiosity. To the idea that anything can be done.

To my sisters, Marie, Stephanie and Chiara. For the support and faith that you always had in me. To the idea that the small things are what matters.

To my uncle and aunt, Eduardo and Rosalba. To my texan cousins, Rosalba, Andrea, Eduardo and Lorena. Always there. There is no way I could have made it to the end of the thesis without you. I will be forever grateful.

You always believed in me. And as long as you believe in me, I can accomplish anything.

ACKNOWLEDGEMENTS

I would like to thank Dr. Ergun Akleman, my committee chair, for providing me continuous support through the numerous months that it took to complete this thesis. I am grateful for his patience and his insight which were both absolutely essential for me to get to the end of this. I would also like to thank the rest of my committee, professor Richard Davison, Dr Ann McNamara and Dr John Keyser for their support during the development of my thesis and also during the numerous semesters at the Viz Lab.

A special thanks to professor Tim McLaughlin for his time and patience, for the direction that he has given to the department and for being a shadow figure in the background of many of the projects and studies I have undertaken in my years at the vizlab.

Thanks also to all the faculty and staff of the department of Visualization for all your time and everything I have been able to pick up from that place, in class, in the lab and even just in the hallways.

I would like to thank the students in the Vizlab for all their support and for everything I have learned by supporting them. The list is quite long and I will include a small selection here: Megha Davalath, Adan Pena, Bobby Huebel, Leticia Silva, Nathan Bajandas, Seth Schwartz, Alan Blevins, Garrett Moring, Ariel Chisholm, Bob Graf, Jon Simpson, Ganesh Rao, Christine Liu and many many others. It was an interesting long long time. There are many more students and professors that have directly or indirectly provided me with challenges that gave me the insight and vision necessary for this thesis.

A special series of thank you goes to my family. All and every single one of them.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE Page

1. INTRODUCTION

The properties of light are extremely complex and difficult to reproduce in computer graphics, but in recent years a considerable number of studies have attempted to accurately model the physical behavior of light through volumes and on surfaces, developing algorithms and techniques that would allow the generation of photorealistic computer generated imagery with ever increasing detail.

Most three dimensional representation models stress photorealism and aim to achieve a high degree of complexity and recreate every possible physical detail. Although all representations are in fact approximations and complete photo-realistic reproduction is unachievable, these attempts continue nonetheless.

Reality is defined by its imperfection and complexity, characteristics virtually impossible to recreate with computers. This intrinsic flaw of computer generated imagery is well known and accepted. While it does not particularly affect the goals of the thesis, it does offer an opportunity.

Similar to more traditional "analogue" artistic models, this thesis aims to offer intuitive control to artists and allow artists to achieve visual results which the audience would consider appealing and believable. The thesis describes an alternative simplified light transfer model not grounded in physically accurate descriptions of the behavior of light, but focused on the aesthetic value and compelling details of the imagery, aiming to achieve an artistic result that the audience would consider believable.

According to Ayn Rand [1], all human directed/generated imagery attempts to offer a visual abstraction. Essential and unique characteristics are selected and integrated into a single model. When referring to an apple and a painted representation of an apple, Rand goes on to note, "No one can perceive literally and indiscriminately every accidental, inconsequential detail of every apple he happens to see; everyone perceives and remembers only some aspects, which are not necessarily the essential ones; most people carry in mind a vaguely approximate image of an apple's appearance. The painting concretizes that image by means of visual essentials, which most men have not focused on or identified, but recognize at once" [1]. The artist performs a process of selective representation and then use different techniques to render the subject of the representation.

The decision to focus this thesis on the visual result of light transport through participating media and the control of such effects is primarily a personal decision. There have been several studies in computer graphics attempting to improve the accuracy of the recreation of the physical behavior of light, but the focus of this thesis is to recreate the perceptual aspects and visual essentials that will allow the viewer to recognize the subject of the representation.

Light transfer in participating media can offer a huge variety of perceptual details, some very evident, such as candle wax or human skin, others a bit more subtle, such as plastic and wood. These details are the result of the high complexity that materials have at a microscopic level and break completely with the common perfection and coldness so common in computer graphics. While highly complex in nature, these details are visual essentials that the viewer recognizes at once and are integral to recreate the believability of the representation.

For example light traveling through the wax of a candle is a visual aspect that is perceptually essential to the image of a candle's appearance, as seen in figure 1.1. Any candle without this visual characteristic would appear unnatural and would clash against the viewer's system of expectations. However most viewers do not have a thorough knowledge of the physical process that is causing these visual details in

Figure 1.1: Photograph of a burning candle.

the wax. The viewer is content considering the simple knowledge that light shines through the wax, without any further explanation or understanding.

A lively glow and soft hue gradients are the results of light transfer in participating media and they are associated with a very natural and often organic characteristic of the objects. In many situations these effects are very apparent, as in the previous candle example or with skin. In these and other situations they are very subtle effects. In all these situations the effect of participating media add a considerable visual interest to the image, offering a sense of three dimensionality and hue temperature of the surface. Further, there is a breaking away from the traditional limitations of simple computer generated images.

The primary and most important difference between a light transfer method based on perceptual goals and one based on goals of physical accuracy of the behavior of light is the knowledge of the subject. The focus of this thesis is to move from a physically based method that is only aware of material attributes to a method that is based on the visual perception experience of the artist.

Perception-oriented visuals are rarely considered in computer graphics. It is important to not mistake non-photorealistic representation (NPR) techniques with perception-oriented, while most NPR techniques attempt to recreate an artistic visual style, these techniques are generally automatic and do not take into consideration the nature of the objects depicted in the scene.

1.1 Artistic Inspiration

There are several rules and features that can define an artistically stunning image. The primary interest of my research is concerned with color and temperature. My goal is to achieve an artistic control and freedom over the hues and values of the image, blending and smoothing the glow of objects and light in the result. Most traditional artistic forms produce perception-oriented visuals based on the artistÕs experience and understanding of the subject.

The perception-oriented goals of this thesis were inspired in a very particular way by the artistic production of Michelangelo Merisi (1571 - 1610), most commonly known simply as Caravaggio.

Caravaggio is particularly renowned for his use of chiaroscuro, a technique that uses light and dark to achieve a definition of volume and form. Caravaggio breaks away from the tradition of symmetrical figures and detailed backgrounds. His work is characterized by a great naturalism and raw realism. In his paintings, the subjects are depicted with an amazing volume and presence, sculpted in a glowing light that offers strong areas of light and areas of darkness.

The Basket of Fruit (1596), seen here in figure 1.2, was painted by Caravaggio in

Figure 1.2: Basket of fruit - Caravaggio, ca 1596.

his earlier years, shortly after leaving Giuseppe Cesari under whom he worked for a couple of years. In this painting Caravaggio is yet to begin using the immense darkness for which he is famous, however Caravaggio's physical particularity in clearly visible in the contrast between the values and tones on the fruit. The result is a strong sense of realism.

The details in the Basket of Fruit are remarkable; Caravaggio's skills are evident and it is easy to observe how he masterfully captures the light traveling through the grapes and the material glowing in the light. Caravaggio's depiction of the

(A) (B) (C)

Figure 1.3: Details from Basket of fruit - Caravaggio, ca 1599.

fruit is not photo-realistic, instead the artist managed to recreate a very convincing representation based on observation and his visual experience of the subject. The attention to detail and the skill of the painter can be seen in figure 1.3.

It is clear that the artist has a very strong understanding of shape, form and light, but this knowledge is not a scientific knowledge, Caravaggio is not attempting to recreate the details of photons bouncing and scattering inside of a grape, but rather he is able to capture the organic perception of light and life in his work.

For years before he painted the Basket of Fruit, while performing support work for Giuseppe Cesari, he had already been studying fruits and other subjects, and through observation and perceptual experience Caravaggio had mastered those visual essentials that a viewer recognizes in a subject such as an organic still life. In figure 1.4 we can observe the visual experience and the painting mastery that the artist is developing in Boy Peeling a Fruit (1593), Young Sick Bacchus (1593) and Boy with a Basket of Fruit (1954) .

The quality of Caravaggio's work increases throughout his career and his organic understanding of light becomes one of the primary characteristics for which he is

(A) (B) (C)

Figure 1.4: Series of early Caravaggio paintings - (A) Boy Peeling a Fruit, ca 1593. (B) Young Sick Bacchus, ca 1593. (C) Boy with a Basket of Fruit, ca 1594.

historically renown.

Caravaggio is famous for his tenebrist style and dramatic lighting. It is possible to observe incredible detail in the glow and hue of his paintings and see how there is clearly a dominant light source, yet the artist organically captures the perceived visual results of indirect lighting. Light is bouncing and wrapping around the forms in of the objects, changing in temperature and hue as it flows thru and around the surfaces. However Caravaggio's forms maintain a traditional monumentality and only in his later work does his form appear less plastic.

The masterful quality of light in Caravaggio's later paintings is evident, but while his painting show incredible control over the aspects of light and volume, the artist has yet to master the effects of light traveling through and around the surface of skin. His lighting on his characters is breath-taking and inspiring, but he is not able to achieve the same level of incredible essential visual elements on his characters' skins as he is able to achieve in his still lifes.

Caravaggio's work is revolutionary in many ways and the goals he achieved influenced a wide range of painters as diverse as Georges de La Tour and members of the Utrecht School such as Gerrit van Honthorst.

Figure 1.5: Supper Party - Gerrit van Honthorst, ca 1619.

Honthorst (1592 - 1656) is considered one of the Dutch Caravaggisti and had spent a short period in Rome learning directly from Caravaggio's work and embracing the italian painter's chiaroscuro technique and mastering the technique in his paintings, often depicting scenes illuminated by a single candle. Not many years after Caravaggio, Hornthorst clearly demonstrates that he has learned from the paintings of the italian master and follows in his footsteps, dominating the perception of light and the interaction between light and volumes. We can admire Honrthorst strive to recreate the perception of light in his painting "Supper Party", seen here in figure 1.5.

Every single person has a strong perceptual experience of what are the characteristics of human skin and how it reacts under different lighting situations, yet even with his amazing and organic results, Honrthorst, just as Caravaggio before him,was not able to completely master these visual effects and recreate their perception, but the attempts to recreate those visual elements were about to be attained.

Figure 1.6: Detail from Saint Joseph - George de La Tour, ca 1642.

Georges de La Tour (1593 - 1652), a contemporary of Horthorst, was able to replicate successfully the perceptual essence of subsurface scattering in painting, with such detail and quality that most advanced physically-based techniques for computer generated images have yet to successfully replicate. La Tour's work is clearly influenced by Caravaggio, but more probably via the Dutch Caravaggisti. He was able to develop the light effect much further than his predecessors had done. In his painting the viewer can perceive the softness and liveliness of the skin as light touches it, with such an organic perception that the viewer can distinguish between the adult skin, child skin and newborn skin. A clear example of his skill and control can be admired in his painting "Saint Joseph", seen here in figure 1.6], completed just a couple of decades after Caravaggio.

1.2 Visual Inspiration

Since the first moments after birth all humans start accumulating perceptual experience and gathering information of how they perceive and recognize the world around them. Everyone has experienced hundreds of thousands of occurrences of the many visual results of light transfer in participating media.

The motivation behind the decision to focus on the effects of light transfer in participating media was not simply intellectual, but strongly influence by the awe and extraordinary beauty of these effects.

Figure 1.7: Photograph of sunlight in clouds

The details of light scattering and reflecting as it travels through clouds is fascinating, and such effect can be observed in very simple situations such as a spring morning or even in more particular situations such as from an aircraft. These visual details in clouds are very common and recognizable, part of each individual's personal perceptual experience of clouds, a photographic example is seen here in figure 1.7. If an artist were to attempt to represent clouds without considering at all these visual essential characteristics, this work would create a clash between the viewer's expectations and the artist's representation. The alternating glow and light absorption allows us to recognize the shape and volume of the clouds.

Figure 1.8: Examples of the effect of light transfer in participating media.

Light travels through most any object or material to varying degrees, sometimes with very subtle results that many viewers may not be able to recognize as present, and at other times the visual effect in the participating media is so important that it defines the visual experience.

Everyone recognizes the effect of sunlight shining through the leaves of a tree, through grapes or through clouds. Everyone has admired light bleeding through the wax of a burning or light creeping through the skin of their hand.

For some time now I have admired the visual effects of glow and warmth that are created in subjects as clouds or candles. The sheen on the surface as light travels and scatters through the volume is extremely compelling and it is essential for the organic and natural appearance of many objects. The smooth glows and gradients in the colors are awe-inspiring and is one of those details that while extremely complex to simulate and accurately recreate, it is instantaneous in its perception. The visual effects and details are limitless and pervasive to our environment. In figure 1.8 I've gathered just a small selection of these awe-inspiring details.

While ambitious, the idea of recreating to some extent such visual quality was among the central motivations of this research.

2. RELATED WORK

Light scattering models are fundamental for realistic image synthesis. A great deal of research has gone into developing sophisticated light transport algorithms to approximate and reproduce compelling realistic results. While it is possible to accurately simulate light transfer through participating media by solving the full radiative transfer equation [2], this would result in an extremely complicated model and be computationally far too expensive for general graphics purposes [3].

The phenomena of light propagation through participating media is very easily observed and while not dominant in many different media, it is essential to recreate the more accurate and minute visual details that more basic computer graphics models fail to achieve.

Light propagation in a participating medium can be described by the volume rendering equation. Four different phenomena must be considered when attempting to comprehend the behavior of light as it travels through a participating medium: absorption, scattering, emission and phase function, shown here in figure 2.1. These phenomena can describe in a general way the physical behavior of light and also the various visual effects that result from light propagation. Modeling differently these phenomena and modifying their influence and importance can result in the description of different media.

The phenomenon of absorption will always reduce the amount of radiant energy in any participating medium and the amount of this reduction can be calculated with the distance the light has travelled through the medium coupled with a coefficient of absorption of the medium. The most simple participating medium only absorbs light with no form of scattering.

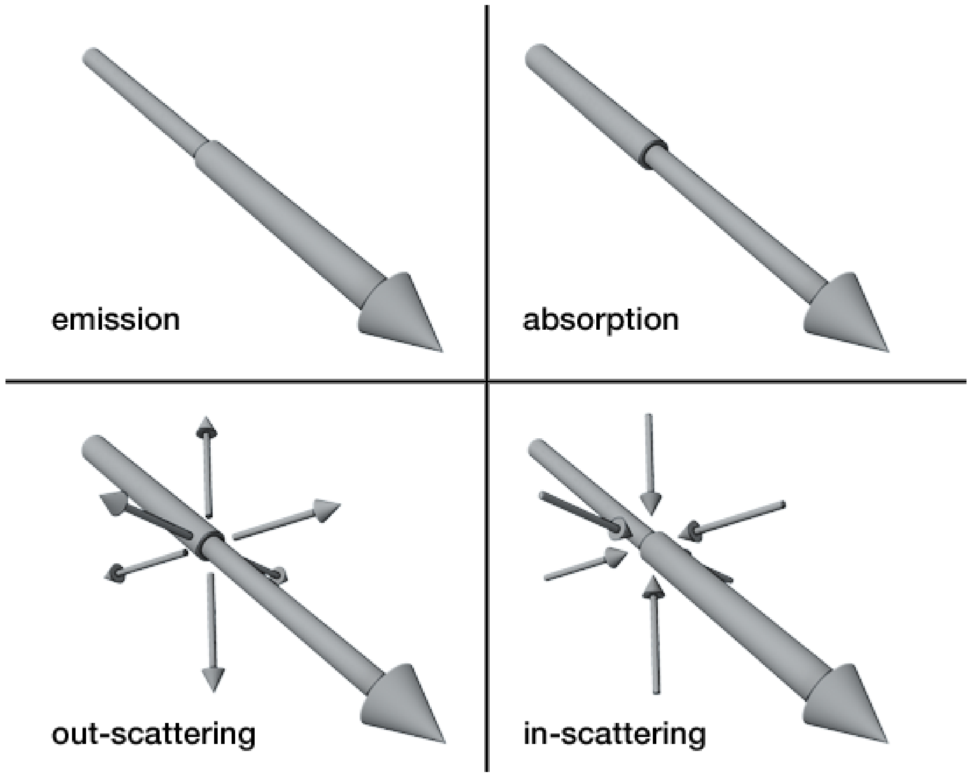

Figure 2.1: Interaction of light in participating media

The scattering phenomenon modifies the direction of the propagation of light. Most participating media can be approximated as a volume of particles having a defined, homogeneous or non homogeneous, distribution. Radiant energy is not just reflected or refracted at the surface points, but also reflected and refracted at every point through the media, scattering in different directions. The influence of scattering can be distinguish in two groups: out-scattering and in scattering. Out-scattering reduces the radiant energy in a particular direction.

The phenomena of emission refers to all the processes that increase the radiant energy in a given direction. Photoluminescence phenomena like fluorescence and phosphorescence contribute to the emission contribution in a full radiative transfer equation but will not be considered in this thesis. The only emission phenomena considered is this model is the contribution of in-scattering radiant energy, as light impinging on a given point in the participating medium that is scattered into the considered direction.

The spatial distribution of the scattered light is model by the phase function of the medium. Different phase functions have been proposed to model different media. The most simple form of phase function is the isotropic phase function, in which light is scattered uniformly in all directions. More complex phase functions are anisotropic and the distribution of the scattered light depends on the angle between the incident direction and the outgoing direction of the light. Anisotropic phase functions allow modeling materials with dominant forward scattering or back scattering.

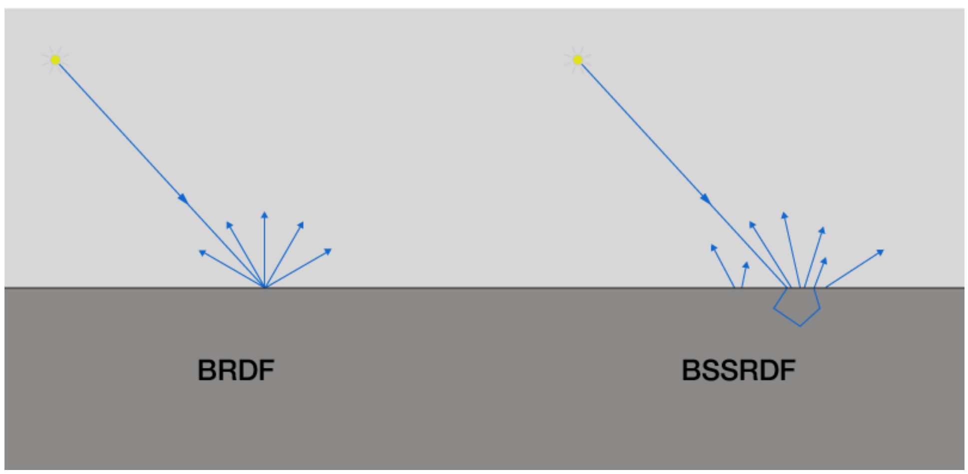

Figure 2.2: no scattering medium vs scattering medium

The most basic models are based on the bidirectional reflectance distribution function (BRDF). The BRDF was first introduced as a simplification of the more general bidirectional subsurface scattering reflectance distribution function (BSSRDF) [4]. While BSSRDF describes the light transport between any two rays that hit a surface, BRDF presents a simplified model that assumes that light enters and leaves the material at the same point. Figure 2.2 visualizes the basic differences between the BRDF model and the BSSRDF model. BRDF creates a very distinct computer-generated appearance because it does not blend any surface features such as geometry, color or shadow. The cold and hard appearance, resulting from this approximation, is valid for metal surfaces, but fails translucent materials, which exhibit many visual surface features due to light transport below the surface.

Most BRDF models approximate surface scattering by a lambertian component, ultimately assuming that light scatters at a single surface point and most models do not model subsurface transport through the material. Subsurface transport can be simulated by solving the full radiative transfer equation, but an accurate simulation would be extremely slow.

Subsurface transport can be simulated accurately but slowly by solving the full radiative transfer equation, and very rarely have papers in graphics taken this approach.

There have been several different approaches to reproduce the effects of subsurface transport, attempting various analytical and numerical solutions to this system. Many of these solutions are organized in the surveys written by Frederic Pérez [5] and then extended further by Eva Cerezo. [6]

Rushmeier introduced the zonal method and extends the treatment received in radiosity from surfaces to volume elements. Rushmeier's method requires computing form-factors between the volume elements and involves a very high order of computational complexity [7].

It is possible to simulate general BSSRDF using numerical techniques such as Monte Carlo path tracing [8][9][10], but as with the previously mentioned methods, these require extensive processing resources and are considerably expensive. Eric P. Lafortune presented a method which extended basic bidirectional path tracing allowing the simulation of basic global illumination effects due to participating media [11]. The resulting image-based algorithm is more versatile, handles multiple scattering in non-homogeneous and anisotropic media, but still is computationally expensive.

In this context Matt Pharr and Pat Hanrahan described a mathematical framework for solving a wide variety of rendering problems and used scattering functions to simulate subsurface scattering results [12].

The photon mapping technique introduced by Henrik Jensen can simulate several of the characteristics of the BSSRDF models, but becomes expensive for highly scattering materials [13][14].

Julie Dorsey simulated full subsurface scattering to capture the appearance of weathered stone [15]. Based on the particle-based composition of stone, this method considers non-homogeneous media and uses a variant of a photon map to capture the in-scattered light in the participating media.

Many of these methods still have difficulties handling refraction while attempting to simulate participating media. Simon Premože formulates an approximate path integral attempting to identify probable paths of light through a medium and efficiently render a variety of scattering materials [16]. Premože also presents a practical method for rendering volumetric effects and approximates multiple scattering effects. The spread of direct illumination is estimated through an expression that captures blurring of radiance and relates scattering and absorption to the distance light travels through the media [17].

Methods that aim to simulate subsurface transport by solving the full radiative transfer equation are capable of reproducing all of the effects of subsurface scattering, but are computationally very expensive compared to the simulation of opaque materials.

Several different methods have been developed to model the effects of subsurface transport, focusing on the amount of scattering in the material (single scattering vs. multiple scattering) and the phase function of different types of materials. Most analytical models of BSSRDF apply separately to two classes of materials. A single scatter material, where light scatters at a single point and only once, offers an analytical solution [18]. These models produce highly directional effects and assume the scattered light propagates from the beams of light propagating in the scene.

Highly scattering materials are solved with different models, often a variation of superposition of a single scattering term and a diffuse term [3].

Jos Stam introduced the use of diffusion theory to approximate multiple scattering at a low cost. He solved a diffusion equation approximation using a multigrid method and achieving a low-cost solution for scenes with optically thick and highly scattering media [19]. Diffusion approximation lowers considerably computation and is able to reproduce the scattering effects in materials with negligible low-order scattering [20], but this method also assumes that phase function of the medium is an isotropic function. This approximation produces incorrect results when the low-order scatter is significant and the phase function cannot be simplified to isotropic scattering.

Other methods attempt to combine the simplicity of diffusion methods with the accuracy of numerical techniques. Several methods use multiple passes to evaluate the subsurface transport and reproduce the effects of subsurface scattering. Henrik Jensen presented a two-pass rendering method for translucent materials that would decouple the computation of incident illumination from the evaluation of the BSSRDF [21] . The first pass would compute the irradiance at the surface, while the second pass would evaluate an approximate diffusion using the computed irradiance from the first pass.

The various approaches aim to achieve an increasing accuracy and lower order of computation to achieve compelling and accurate results. Jeppe Frisvad introduces a method for accurately computing the scattering effects in a participating medium with accurately defined physical and optical components. This method is based on a robust generalization of Lorenz-Mie theory of scattering radiation [22]. Wojciech Jarosz presented a novel radiance caching method that allows an efficient rendering of participating media using Monte Carlo raytracing [23]. This method can reproduce light scattering in both homogeneous and non-homogeneous materials. Raanan Fattal presented a method which models light absorption and scattering using propagation maps which offered high-quality approximations of several different media [24]. Craig Donner proposed a method that would allow considerable speedup while recreating medium that cannot be approximated as a single scatter medium or a multiple scatter media [25]. This method allows the representation of complex spatial and directional dependent media that have anisotropic phase functions.

3. METHODOLOGY

This thesis presents a an alternative approach for light transport based on perception and artistic interest. Unlike physically-based models that attempt to accurately simulate the behavior of light as it travels through a volume, the approach we present in this thesis simplifies the behavior of light and values visual interest above physical accuracy.

The most important feature of our approach is its limited complexity and the direct correlation between the visualized data and the final result. The model is intended to be intuitive to the user and allow them to modify and control the illumination data set and obtain predictable results.

The approach we propose in this thesis can be structured into two interacting components which allow the artist to achieve his visual goals. The first component is the rendering method, which is the technical core of the approach and defines how the illumination data is computed to deliver a final image. The second component of this approach is the conceptual core of the thesis. Our primary goal is the painterly approach and intuitive approach to artistic control of the illumination data and how this data will define the visual output.

3.1 Rendering Method

This approach introduces a simplified model that takes into account exclusively diffuse scattering and an isotropic phase-function and identifies the uniformly distributed scatter as the dominant component of the visual results of light transfer. This simplification allows the model to perform multiple elaborations of the illumination data independent from the environment illumination. This assumption allows us to not consider other forms of scattering contribution such as specular scattering or scattering effects typical of anisotropic phase functions.

Based on this model, the thesis presents a rendering method that consist in 4 stages:

- discretization
- direct illumination
- indirect illumination - scattering
- final gather stage

3.1.1 Discretization

The first stage of the rendering method transforms the polygonal representation into a three dimensional volume representation in which it will be possible to approximate the behavior of light transfer. The discrete approximation of the volume allows us to simplify the computational cost of this approach and also maintain a more intuitive model to the artist. While by far not the only option, in this method we have decided to to use a voxel approach to the volume representation.

Three different representation forms for a circle are illustrated in figure 3.1.

A voxel (also known as volumetric pixel or volumetric picture element) is a volume element, representing a value on a regular grid in three dimensional space. This is analogous to a pixel, which represents 2D image data in a bitmap. The position of a voxel is inferred based upon its position relative to other voxels. While polygons are able to efficiently represent simple 3d structures with large amounts of empty or homogeneously filled space, voxels can easily represent regularly sampled spaces that are not homogeneously filled. A voxel represents a single sample, or data point, on a regularly-spaced, three dimensional grid. This data point can consist of a single piece of data or multiple data. A voxel intrinsically represents only a single point on

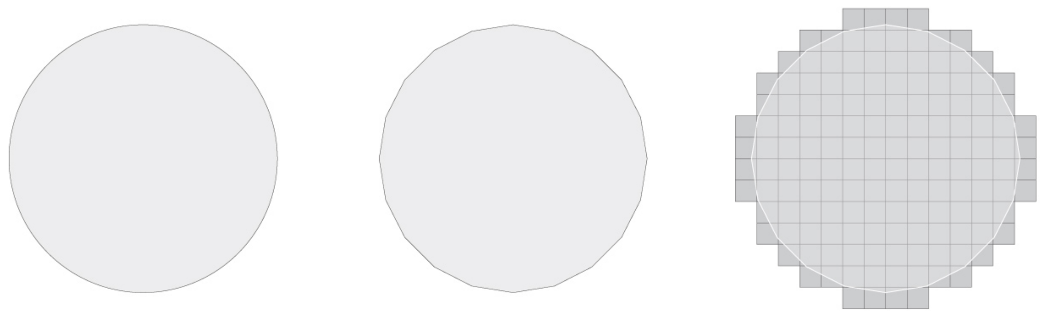

(A) (B) (C)

Figure 3.1: Different representations of the same circle object - (A) Parametric representation (B) Polygonal rappresentation (C) Discrete voxel rappresentation

this grid, not a volume; the space between each volume is not represented in a voxel-based dataset. This missing information may be reconstructed or approximated, e.g. via interpolation. Our approach proposes a discrete approximation of the volume aiming to simplify the difference. In figure 3.2 we can observe the difference between the polygonal representation of an object and the voxel representation of the same object.

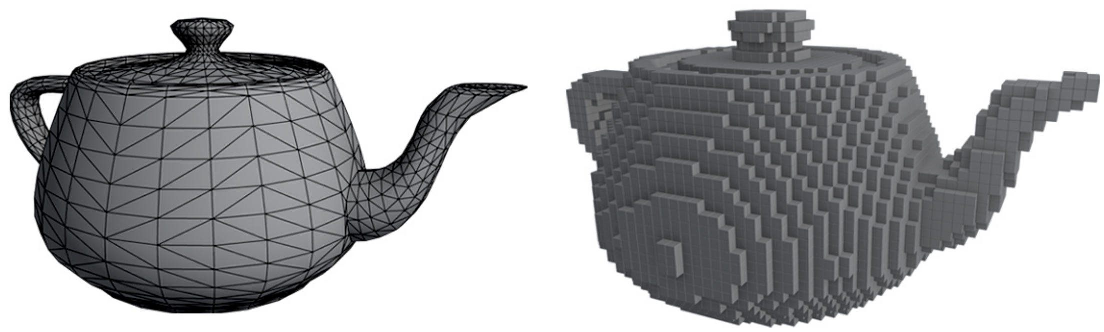

Figure 3.2: Mesh teapot and Voxel representation of the same model

The volume will be represented by a voxel-set defining a given point contained inside the polygonal hull of the original surface. The different values in the volume are approximate via interpolation between the values of the closest voxel points.

The resolution of the voxel set can have considerable influence on the final rendered image. More accurate and believable images will result from a voxel resolution that would limit the amount of change of the surface normal throughout the voxel.

3.1.2 Direct Illumination

Once the object has been discretized and each object is represented by a surface and a corresponding volume element set, it is possible to calculate the contribution each light in the environment has on each single volume element.

At this step our focus is light attenuation due to the participating medium. As light is traveling through the medium a combination of absorption and scattering cause a decay in its intensity. Figure 3.3 visually represents the relationship between the distance and the amount of light absorb by the material.

In optics, the Beer-Lambert law relates the absorption of light to the properties of the material through which the light is traveling [26].

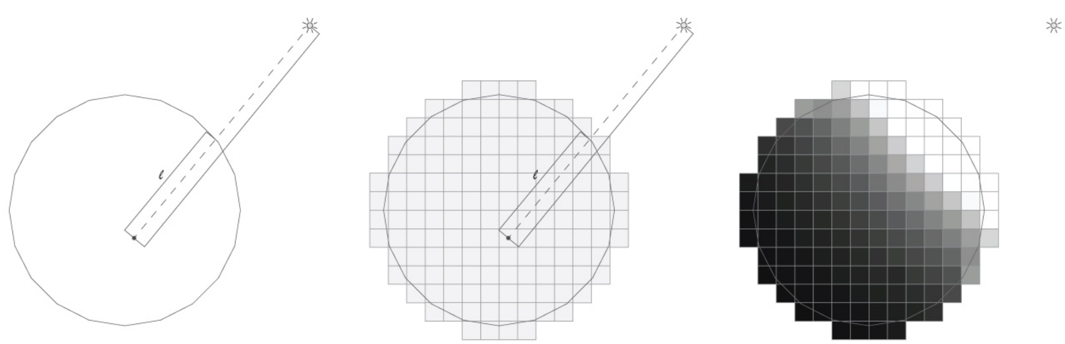

(A) (B) (C)

Figure 3.3: Approximation of direct illumination - (A) Calculate distance from light source and the path length through the media, for a given discrete point (B) Voxel representation (C) Approximated illumination values for each voxel

The law states that there is a logarithmic dependence between the transmission of light through a substance and the product of the attenuation coefficient of the substance, α, and the distance the light travels through the material (i.e., the path length), ℓ.

We can determine the measured intensity I of light transmitted through a layer of material with thickness, ℓ, related to the incident intensity I_0 according to the inverse exponential power law:

$$I = I_0 e^{-\alpha\ell}$$

where ℓ is the path length [27].

This allows us to determine the direct illumination of each volume element that describes the object. The only parameters that are necessary to determine the light intensity at the single volume element is the attenuation coefficient of the medium and the path length of light traveling through the medium to reach the volume element. In Figure 3.4 we can observe an example of the direct illumination through the volume of a voxel representation of a teapot and a section which allows us to observe the decay through the volume.

The process of calculation of the direct illumination of each volume element is therefore independent from all other volume elements representing the object.

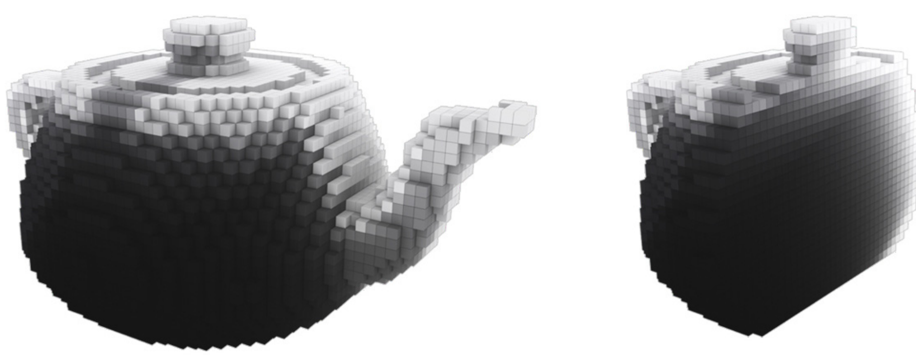

Figure 3.4: Voxel representation of direct illumination from point source and cross section of representation

The intensity of the direct illumination, determined by the attenuation coefficient

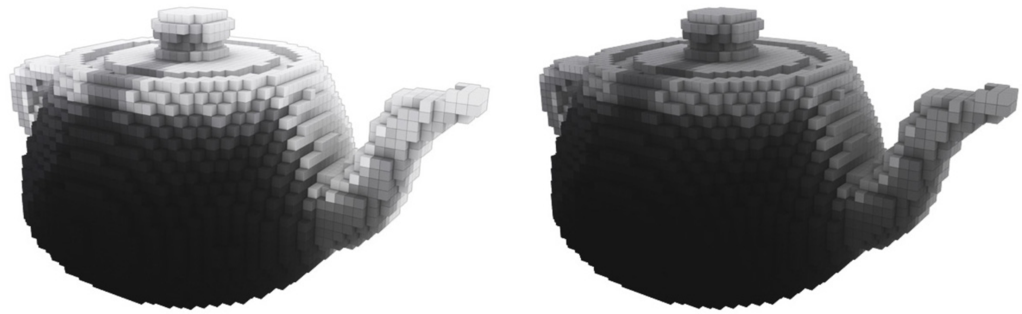

(A) (B)

Figure 3.5: Direct and scattered illumination values per voxels - (A) Illumination values of the light reaching each voxel (direct) (B) Illumination values of the light being scattered at each voxel

of the material, allows us also to determine how much of the attenuated radiant energy was absorbed and how much of this energy was scattered. This will allow us to complete this stage with two separate voxel-sets, one defining the direct illumination for each volume element and one defining how much light is scattered at each volume element. Figure 3.5 shows the two output voxel-sets from the direct illumination stage.

3.1.3 Indirect Illumination - Scattering

The third step of our approach begins the approximation of the scattering contribution through the volume. Once the primary direct illumination and scattering radiant energy were determined by the previous step, it is possible to compute the contribution of the scattered light inside the volume by calculating the scattering influence received at each volume element from the other volume elements of the volume.

Our approach assumes an isotropic scattering in the medium and this simplifies the calculation necessary to determine the contribution of scattered illumination. This step is completely independent of any light setup in the in the scene and only

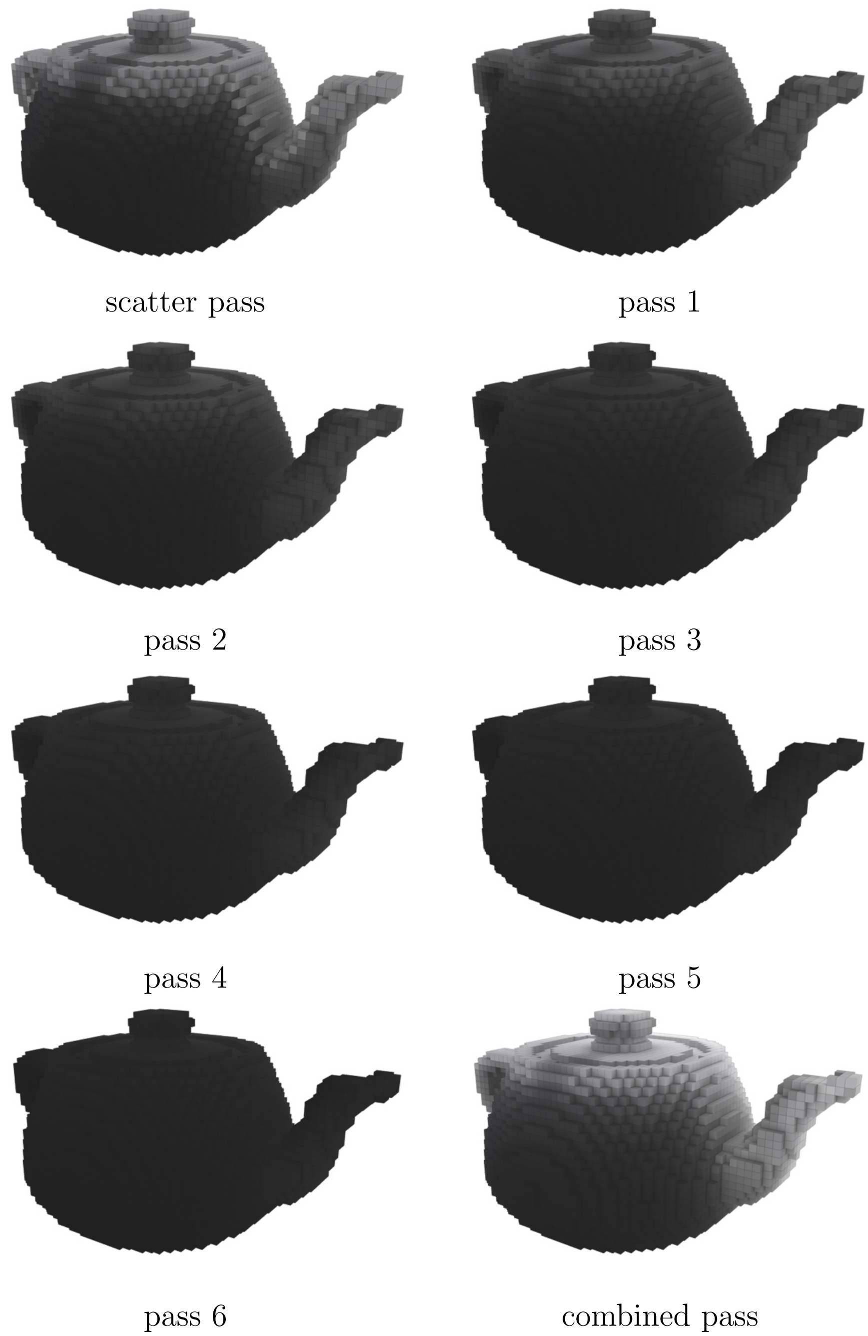

Figure 3.6: The effect of light transfer in participating media - teapot object

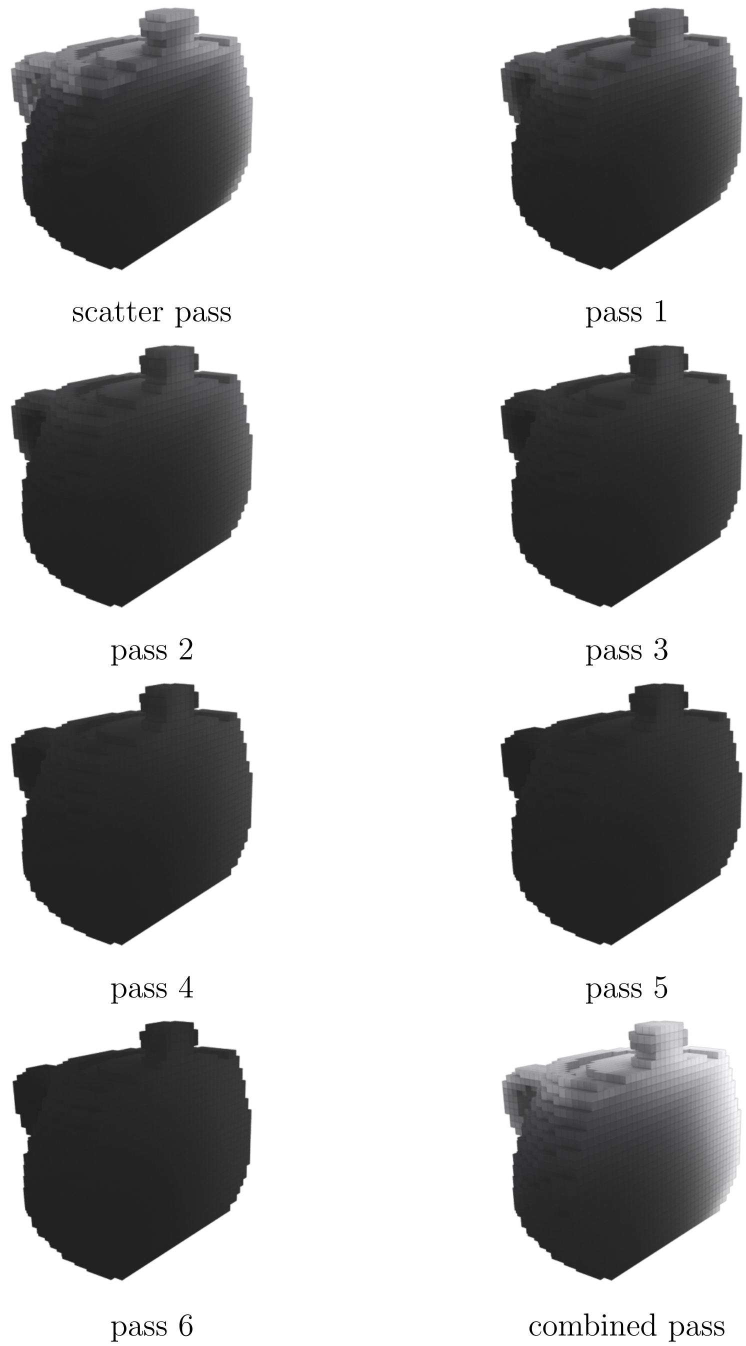

Figure 3.7: The effect of light transfer in participating media - cross section of teapot object

takes into consideration the values stored in the voxel-sets defined in the previous step. In this step we trace the contribution of every volume element in the voxel-set to each single volume element, approximating the first scattering contribution through the volume. We proceed to gather this information and describe a new voxel-set. This step is repeated multiple times to reproduce the continuos scattering through the material. Figure 3.6 and Figure 3.7 show the visual result of each pass and the result of the accumulated scatter.

The influence of the scattered light from the voxel-set to each single volume element still follows the Beer-Lambert law.

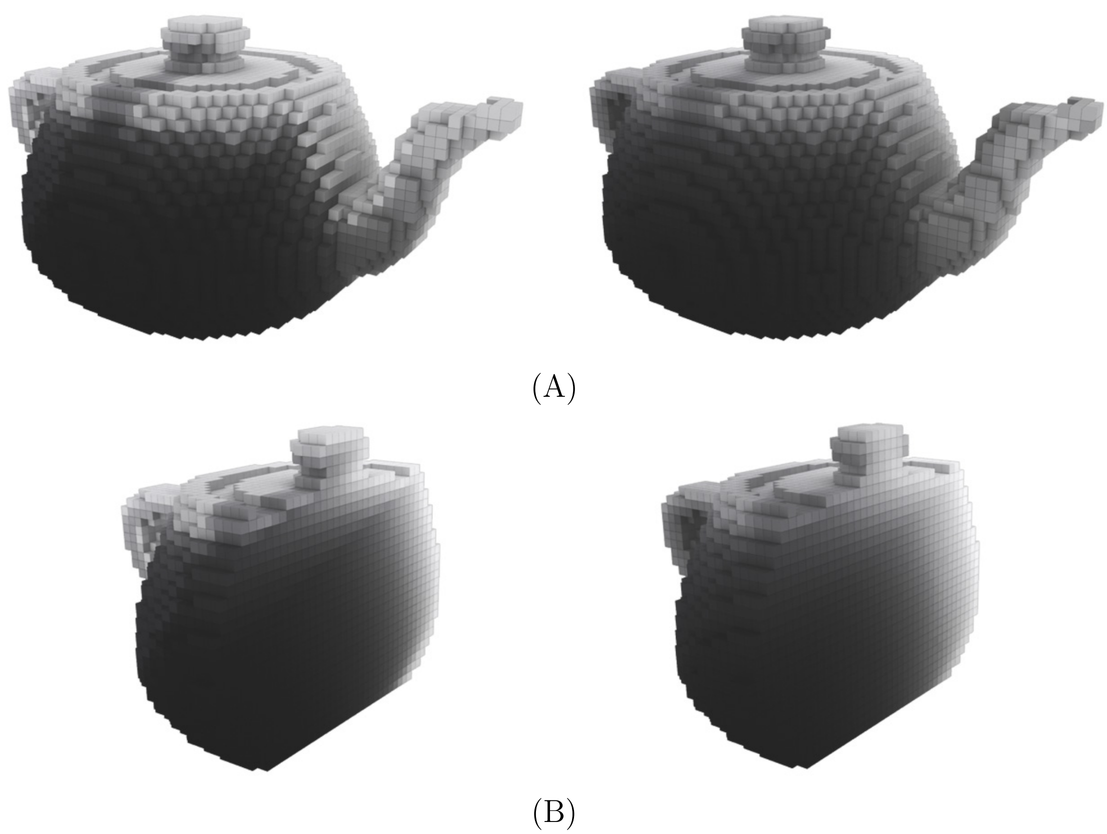

Figure 3.8: Comparison of the illumination values of direct illumination and multi-pass scatter illumination.

By repeating this step multiple times using the voxel-set defined during the previous pass, we approximate the effect of multiple scattering. In Figure 3.8 we can observe how the multiple passes recreate the smoothness and glow common to many subsurface scattering effects.

3.1.4 Final Gather stage

The final step of this approach gathers all the illumination data from the various voxel-sets and defines a final discrete volumetric representation of all the illumination, direct and indirect. Once this final voxel-set is defined the final step of our approach uses this data to determine the subsurface scattering contribution at each point on the original surface.

We used a final gathering technique to estimate the influence of the voxel-set on the surface. Traditionally final gathering is a technique for estimating global illumination for a given point by either sampling a number of directions in the hemisphere over that point (such a sample set is called a final gather point). In this final stage we used the same principal by sampling a number of directions in the volume from the point.

3.2 Artistic Control

The artistic control component of this approach is the core element of this research. The control component is not separate from the rendering method, but rather they are tightly intertwined to achieve the painterly approach. The rendering method allows a couple of moments in which an artist can influence or control the process that will generate the final visual result.

Considering the structure of the rendering method, there are two different ways the artist can direct the visual results of surface scattering.

The first way the artist can influence the final image is by modifying the lighting setup during the direct illumination stage of the rendering method. The rendering method allows to use an independent lighting setup unrelated to the lighting setup used to compute other lighting contributions (diffuse, specular, reflections etc). Any lighting setup used during the first illumination data set, therefor by modifying the lighting setup the direct illumination stage will describe a different voxel set and consequentially the rendering method will generate a different visual result. This allows the artist to achieve his desired look, however the correlation between the lighting setup (positions, colors and intensities of the various lights) and the visual effects of light scattering through participating media is not very intuitive and sometimes can require extensive trial and error.

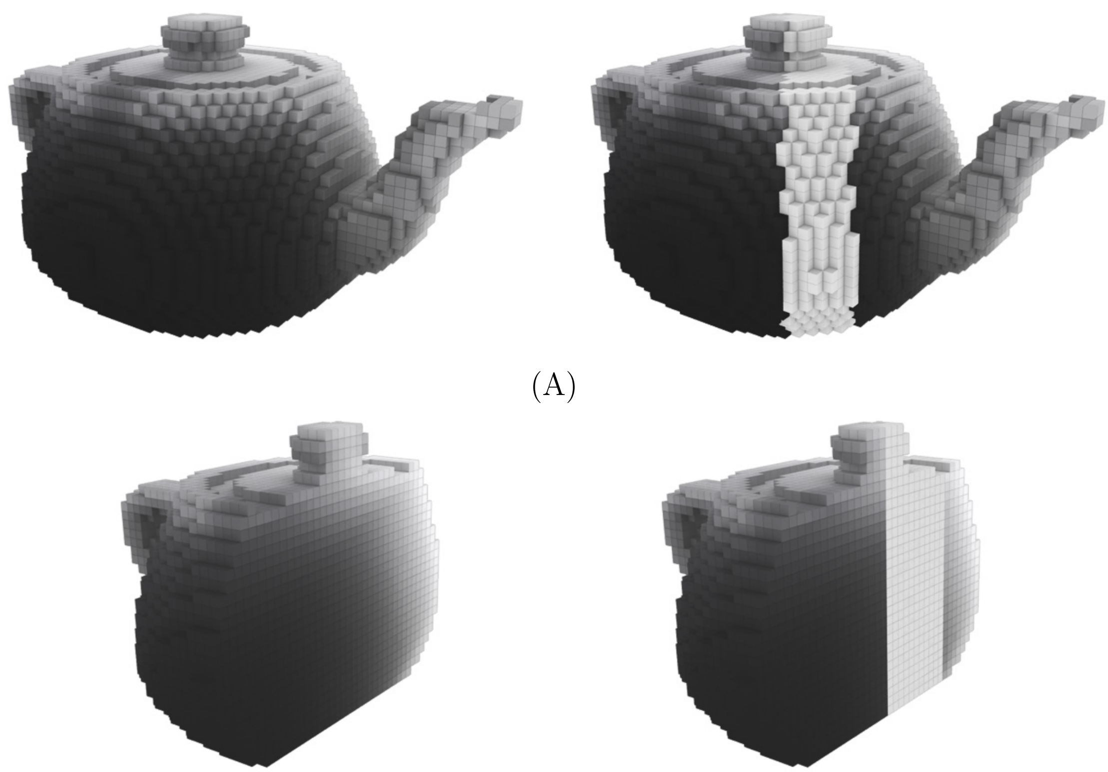

(A)

(B)

Figure 3.9: User controlled edits to the values of the voxel-set

The second opportunity of control lies in the volumetric illumination data. The voxel-set, result from the direct illumination stage of the rendering method, is the central element of our approach. One can consider the rendering method as a process that generates a voxel-set which approximates the illumination data of the volume, followed by successive stages which build on and modify various voxel-sets until a final approximation of all the volumetric illumination data is achieved. Finally this final voxel-set defines the contribution of light transport through the volume.

The goal of this thesis is to offer an intuitive correlation between the volumetric illumination data and the final visual results of light transport through participating media. The discrete representation of the volume offers the artist this correlation. By offering to the artist the opportunity to modify and control values and parameters of the illumination data set, our approach allows the artist to "paint" the effects of scattering to the final image. Figure 3.9 is a very basic example of user edited values in a voxel-set.

It is possible to tweak and modify the voxel-set and the illumination values throughout the indirect illumination stage of the rendering method, therefore generating an art-directed final illumination data set.

The correlation between data set and the final result is considerably more intuitive than the first control method.

4. IMPLEMENTATION

The painterly approach we proposed in this thesis was implemented in a predominantly modular way, attempting to compartmentalize the components and the different stages of the process.

The first part of our implementation is to offer a simple framework which will accomplish the rendering method, while the second part of the implementation will focus on the tools and interfaces which offer the artist control over our approach.

The implementation follows closely the concepts and the process stages introduced in the previous chapter.

4.1 Rendering Method

The rendering method is implemented primarily in C++. Each major stage of the method is structured into its own set of tools, while sharing a basic set of methods and classes common to the overall rendering method.

One of the core concepts of our approach to light transfer in participating media is the approximation of a volume in discrete units. Each of these units also stores material parameters and lighting information gathered from the scene.

The implementation is built around a simply defined data structure and simple ascii text document for input, output and archiving discretized volumetric data, which we defined as the VXF file format.

4.1.1 VXF Format

The VXF format is a very simple format developed with the single purpose of storing a list of volumetric elements and several types of data associated with them. The format is composed in two principal sections: a header section and a list section.

The header section contains basic information which defines the voxel-set representation of the volume and its overall set parameters. In the header we store the width, depth and height of the volume's axis-aligned bounding box, along with the size of a single voxel element. Additionally the header stores the minimum and maximum positions of the volume set expressed in the world space coordinate system. These parameters define the dimension and position of the entire voxel-set.

Here we present an example of the structure of a VXF file header.

```
size 64 32 41 1.0

min -30.0 -10.0 -20.0

max 34.0 22.00 21.0
```

The list section of the VXF format contains a list of volumetric elements, each element with a determinate series of parameters specifically defined. This list format allows for a versatile use of the VXF format, considering that any series of parameters can be stored in each volumetric element. The format uses a sparse organization of the data information and is uncompressed.

The parameter series is not defined in the format; it can potentially store any sequence of parameters for each volumetric element. In our implementation the parameter series is primarily used to store a boolean value, which indicates if the voxel is part of the participating media, and RGB values, which store the illumination values of the voxel.

4.1.2 Discretization

We chose voxelization as the discretization method for our implementation. This method was also used as base method in the definition of the VXF format and other implementation design decisions. The voxelization process is implemented in C++.

The implementation requires as input a triangulated obj mesh and an optional voxel dimension.

The volumetric representation resulting from this process is a dense voxel-set. The output of this process is stored in a VXF file with a single parameter stored in the parameter list, a boolean parameter indicating the inclusion of that voxel in the volumetric representation of the participating media. Figure 4.1 illustrates the the relationship of the obj model and the voxel-set created by the voxelization program.

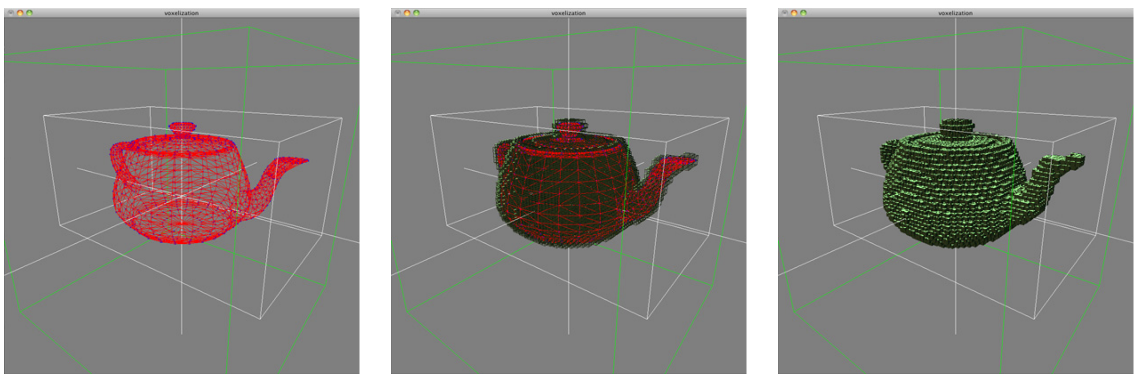

OBJ mesh | Voxels and OBJ mesh | Voxels

Figure 4.1: Process of discretization - Voxel program

4.1.3 Direct Illumination

The direct illumination stage is implemented in C++ and it extends on the output of the voxelization stage. This primary illumination process requires a wider set of input elements and a more extensive calculation.

The direct illumination process require a triangulated obj mesh, a illumination setup, an output VXF file from a voxelization process which defines a volumetric representation, and a series of material parameters. Each voxel is processed independently, which allows this stage to be multithreaded in our implementations.

Different illumination setups have been implemented for this stage. The types of illumination available at this stage are:

- pointlights
- spotlights
- directional lights
- area and volume lights
- environment lights

Given the illumination setup, the luminance contribution is calculated for each voxel. The luminance is calculated by taking into consideration the materials coefficients of absorption, decay, and scattering, the distance light has traveled through the participating media and the basic lighting parameters of the given illumination setup. This step calculates the resulting illumination and takes into consideration shadowing, absorption and other light characteristics as intensity, color, and falloff.

In the implementation we generally use material coefficients constant throughout the entire volume, but the implementation also allows to define different material coefficients throughout the volume. The implementation can also use material coefficients from another VXF file that has coefficients of absorption, decay, and scattering stored in the parameter list for each voxel.

The direct illumination process outputs two separate sets of illumination data, both stored in separate VXF files. These VXF files store in the parameter list of each volumetric element a boolean parameter indicating the inclusion of the voxel in the volume of participating media, and also stores a luminance parameter stored in RGB form. The first of VXF files stores in the these parameters the direct illumination values each voxel receives, the second VXF file stores the amount of scattered luminance emitted from each voxel.

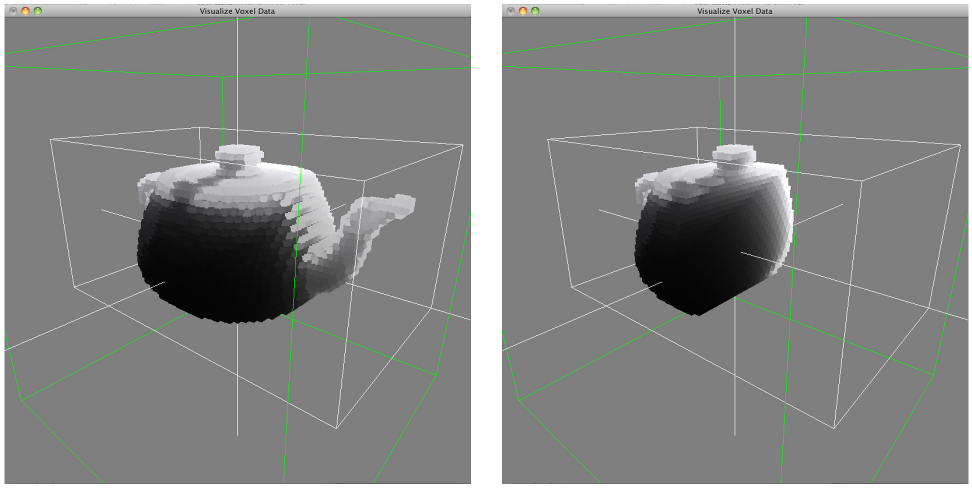

Figure 4.2: Direct illumination from a pointlight source - cross section.

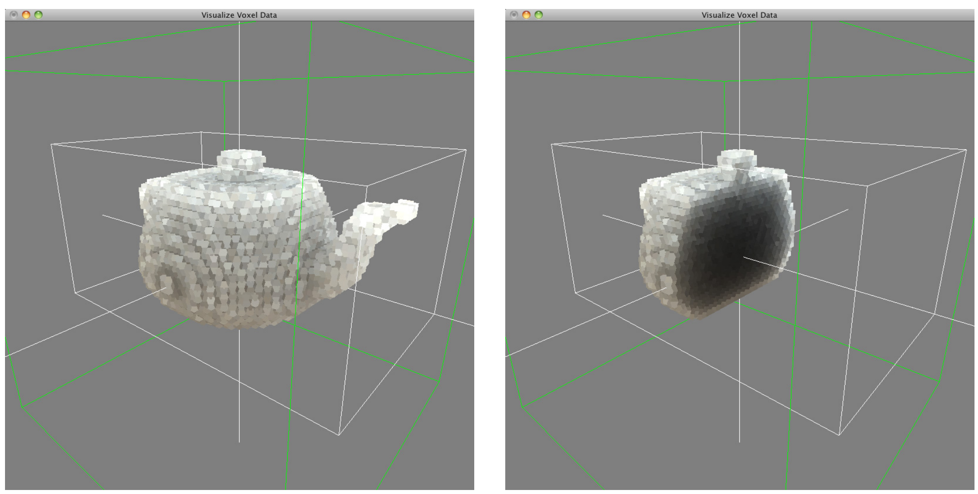

Figure 4.3: Direct illumination from a environment light source - cross section.

The direct illumination stage can become more expensive from a computational standpoint, especially if the illumination setup uses more expensive illumination types such as environment lights or area and volume lights. However the luminance of each voxel is computed independently from all the other voxels and this allows this stage to be multithreaded in our implementation.

It is also important to remember that in case the desire is to integrate the participating media effect into a scene, the illumination setup must take into consideration such scene to have correct shadowing and illumination data.

Figure 4.2 and figure 4.3 show in the implemented application the result of the direct illumination using different light setups.

4.1.4 Indirect Illumination - Scattering

The scattering stage is also implemented in C++ and uses VXF files as primary form of input and output of illumination data. This stage does not require a mesh, the entire calculation is executed on the illumination data stored in the input VXF file.

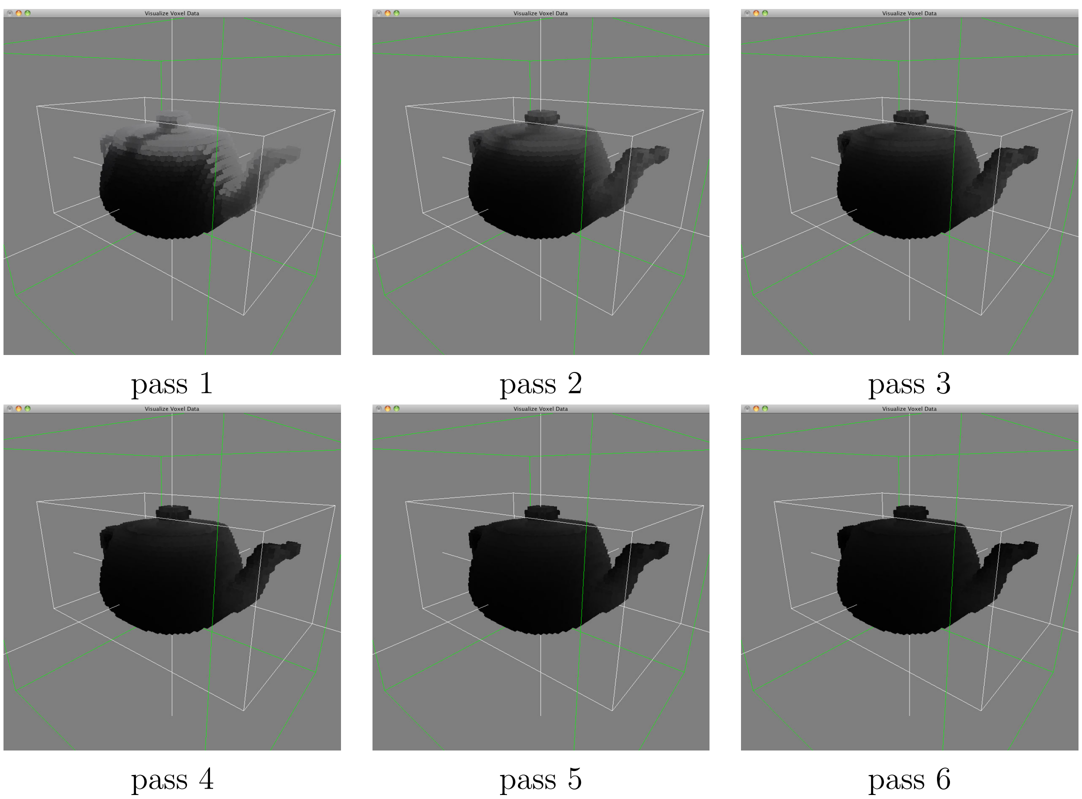

Figure 4.4: Scattering passes based on point light illumination

Figure 4.4 and figure 4.5 show in the implemented application the result of mul-

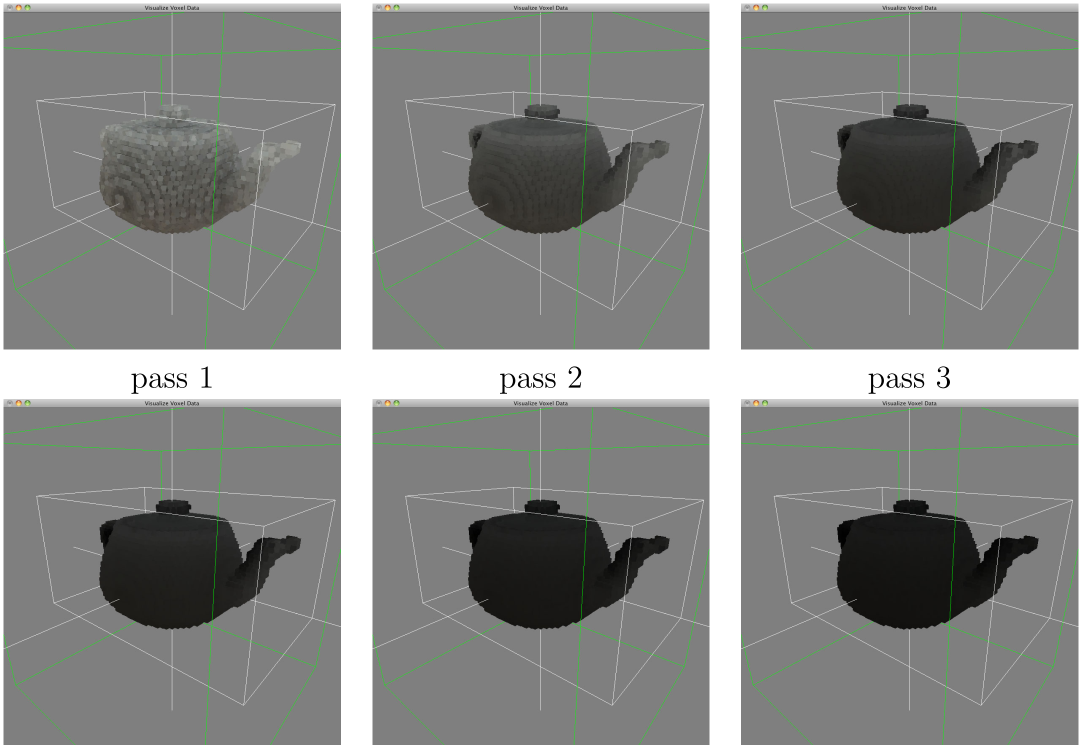

pass 1 pass 2 pass 3

pass 4 pass 5 pass 6

Figure 4.5: Scattering passes based on environment light illumination

tiple scattering passes starting from a voxel-set illuminated with a point light and the same voxel-set illuminated with an environment light.

All the necessary information regarding the illumination setup and the mesh were already elaborated in the previous stages and stored into the illumination data.the material parameters of scattering, absorption, and decay are still required in the stage to determine the output illumination data.

The objective of this stage is to simulate the contribution of scattered light from the entire voxel-set to a single voxel. Several different techniques were tested during our research to determine which technique would give the better result and offer more opportunity for control further along the project.

The technique used in the final implementation is based on the principal of volume ray marching. At each voxel we performed a final gathering pass and sampled in a number of directions from the center of the voxel. The implementation would trace a ray along the direction and march through the volume. Along the path of the ray, equidistant sampling points are selected. In general sampling points will usually be located in between voxels. Because of that, it is necessary to interpolate the values of the samples from its surrounding voxels.

This process is repeated multiple times, using as input the illumination data set resulting from the previous process pass. This makes it possible to simulate multiple scattering and and participating media with higher scattering coefficients.

The multiple scattering pass are then accumulated into a single illumination data set which is stored in an output VXF file. In our implementation it is also possible to output each single scattering pass to a separate VXF file.

4.1.5 Final Gather and Rendering

The final stage of the rendering method is implemented as an extension of a C++ raytracer. All the computation necessary for the contribution of light transfer in participating media is implemented in a shader calculation. This feature is added with a series of extra parameters required by the shader. The participating media shader requires an input VXF file with RGB illumination data stored in its parameter list, material coefficients of absorption, scattering and decay.

The effects of light transfer in participating media are calculated per shading point and at this stage are mainly independent from the scene's light setup and other scene parameters. All the illumination data necessary for our method is already stored in the input VXF file.

The objective is to integrate the contribution of the illumination data in the voxel-set to the single shading point. The technique used in the rendering stage to


point illumination environment illumination

Figure 4.6: Scattering contribution

accumulate the light scattering through the volume is very similar to the technique used in the scattering stage. Starting at the shading point we sample the illumination data in a number of directions. The implementation performs a final gathering pass and along each direction selects equidistant sampling points. It is necessary to interpolate the values of each sample from the surrounding samples.

We can observe in Figure 4.6 the results of the final gathering implementation in the C++ raytracer. The two examples show the visual contribution of scattering resulting from a point light and from an environment light.

Due to the approximation caused by the discrete nature of voxels, it is necessary to add a weighting to each sample and consider invalid ray directions and sampling points that would not be inside of the volume defined by the mesh. This becomes a very common need in situations where the change of curvature on the surface in a single volume is elevated. While weighting and eliminating some samples can offer in many situations acceptable results, in extreme cases this will cause noticeable anomalies and discrepancies.

4.2 Artistic Control

The addition of a user interface that allows artistic control to the rendering method is the core element of the painterly approach of our method.

The first approach the artist can employ to influence the effects of light transport is to directly modify the lighting setup. This approach is very common to most computer graphics techniques, but not as intuitive and effective as our approach aims to offer. While the user can tweak the illumination intensity values and colors, and modify the illumination contribution coefficients, even this approach only offers the user indirect control of the effect of light transport through participating media. In order to offer the user direct control over the volumetric illumination data that is stored in the VXF files we have developed a set of tools deployed in Autodesk Maya and a set of procedures and operations in C++.

4.2.1 VXF interface

The purpose of this interface is to offer the user a realtime tool to visualize and edit the volumetric illumination data stored in a VXF file. The VXF interface is an Autodesk Maya tool developed in python. This tool allows the user to load any VXF file in the Autodesk Maya environment , andvisualize and access the complete volumetric data as a particle set. This allows the user to explore the data in a 3D viewport and take advantage of Autodesk Maya's camera and navigation.

The interface loads the data from a VXF file and organizes the data, offering a series of tools to edit the voxel values independently.

Figure 4.7 shows a screenshot of the Autodesk Maya environment with a visualized voxel-set, while figure 4.8 shows the UI for one of the voxel editing toolsets.

The interface offers the user tools to edit the values of each volumetric element. The user can modify the stored value by hue, saturation and value, and also directly set the RGB values of the selected voxels. It is also possible to add and remove volumetric elements to the set.

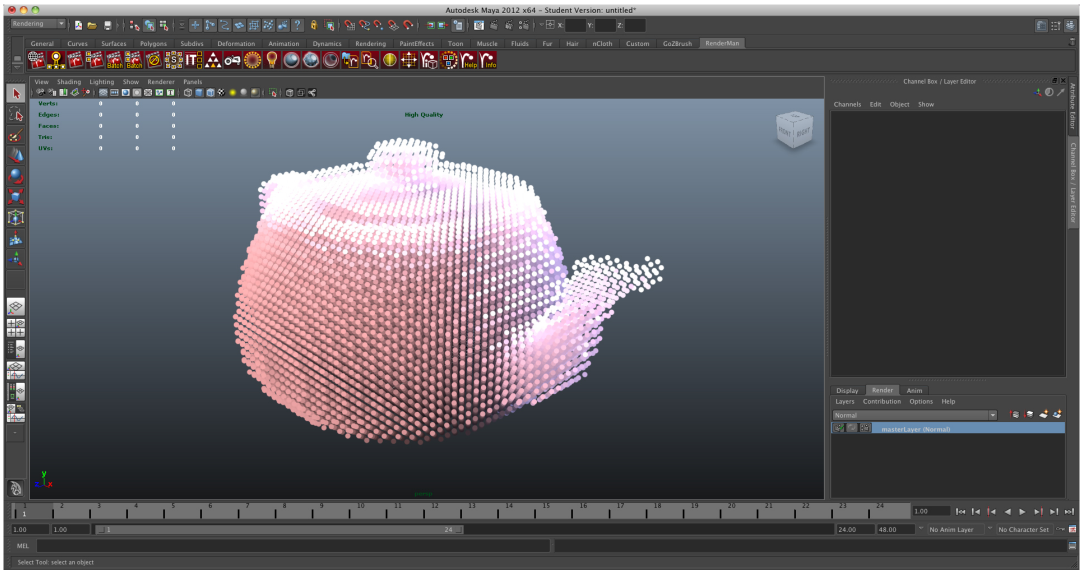

Figure 4.7: Screen capture of an imported VXF data-set to particles in Autodesk Maya

Once the user is satisfied the interface presents a series of different settings to export the edited volumetric data for further uses in the rendering process. The VXF interface offers the following export settings:

- Basic export. This setting exports the entire edited volumetric dataset as a new VXF file.

- Difference export. This setting exports only the changes the user has made to the original dataset. The primary reason the user would choose this setting would be to edit externally the changes made to the original dataset using some of the procedural tools developed in C++.

- Difference + smooth export. Very similar to the Difference export setting, this setting exports the changes the user made to the original dataset and then

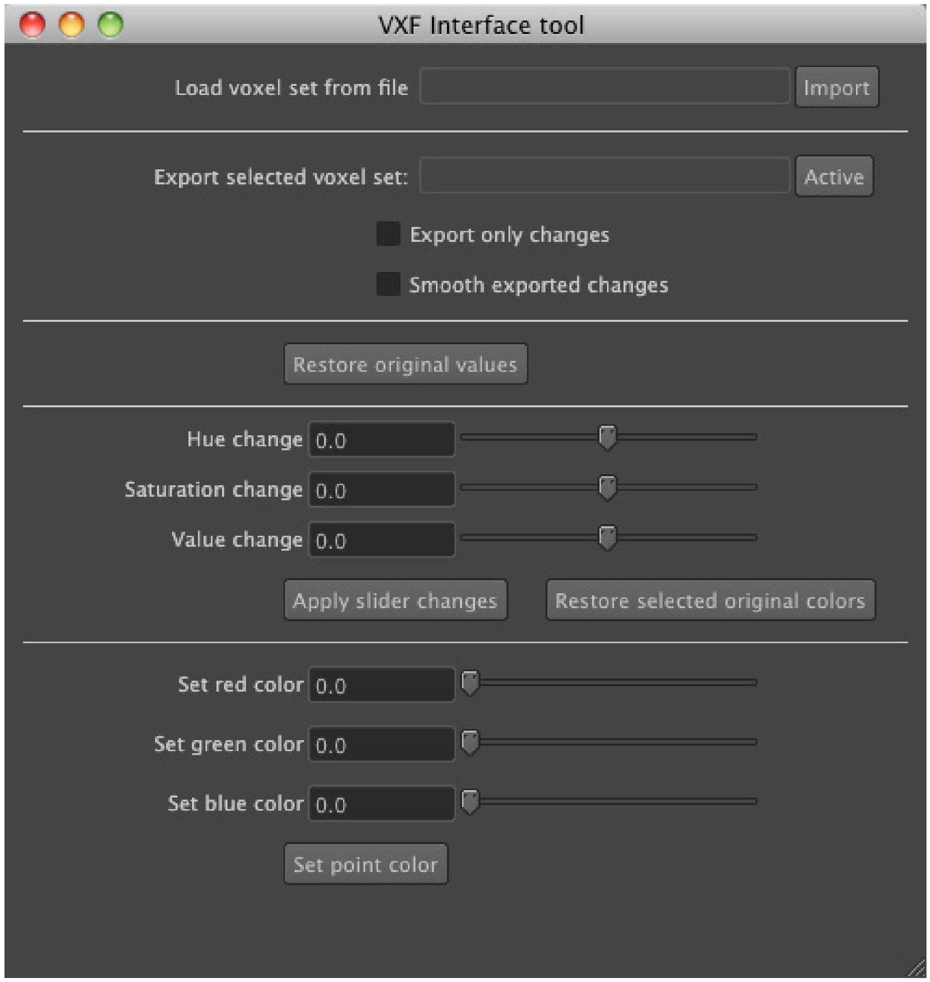

Figure 4.8: VXF Interface Tool

proceeds to procedurally smooths the border voxels before compounding these smoothed values with the original dataset.

4.2.2 VXF procedures and operations

Our research offers an artist control over the visual effects of light transport by giving the artist tools to manipulate the illumination data in the volumetric data structures and stored in the VXF files. While the VXF interface was designed to directly manipulate singular voxel parameters and values, most of the procedures and operations perform a particular computation on each single value of the voxel-set.

The different procedures can be organized based on the type of operations performed and what are the input parameters of these operations.

4.2.2.1 Blending operations

The blending operations use two voxel-sets with equal set parameter (dimension, scale and bounding box location) and perform traditional blending functions between the color values of the respective voxels.

In order to offer the user as much artistic control a possible we have recreated a series of blending operations similar to some of the blending operations offered in common digital photo-manipulation packages.

Considering that in our rendering method the most common values stored in the voxels are the illuminations values, it is acceptable that these values are not contained in a basic RGB spectrum from 0 to 1. Negative values and values greater than one are acceptable and actually relevant considering the desired visual results and nature of the effect this project is attempting to recreate and artistically control the effects of light transport.

We implemented and offer the artist the following short list of blending operations:

- Add

  This operation adds the voxel illumination values of the primary voxel-set and the voxel illumination values of the secondary voxel-set. An example of this operation can be seen in figure 4.9.

  The result of this operation is usually an overall brighter voxel-set, even though this is not always the case since it is possible for the artist to create voxel-sets that contain negative illumination values. This operation is commutative and one of the basic operations extensively used by the VXF interface. This blending operation is also known as Linear Dogde in the Adobe Photoshop Package.

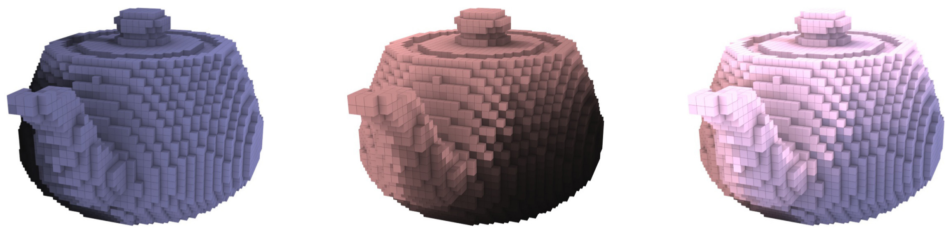

VXF A VXF B VXF A + VXF B

Figure 4.9: Blending Operations - Add

- Subtract

  This operation subtracts the voxel illumination values of the primary voxel-set and the voxel illumination values of the secondary voxel-set. An example of this operation can be seen in figure 4.10.

  This operation is not commutative. This operation usually produces overall darker voxel-sets, but this is not always the case, since it is possible to have voxel-sets with negative illumination values created by the artist.

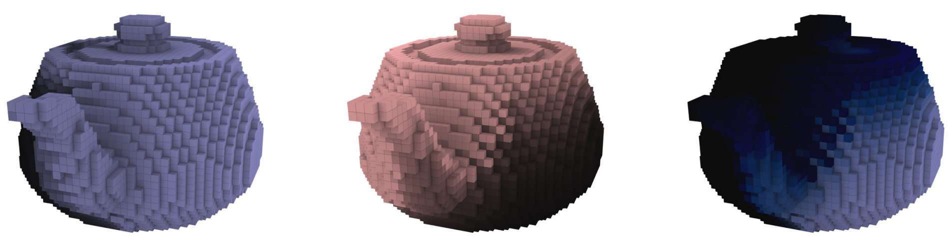

VXF A VXF B VXF A - VXF B

Figure 4.10: Blending Operations - Subtract

- Multiply

The multiplication operation multiplies the voxel illumination values of the primary voxel-set and the voxel illumination values of the secondary voxel-set. An example of this operation can be seen in figure 4.11.

Due to the operation between the two voxel-sets, this operation also offers a flag to reduce all negative values in the voxel-sets to 0 and not generate excessively negative values. It also offers a flag to normalize the illumination values of the secondary voxel-set, allowing a more predictable result. These flags were added after different iterations and attempts of artistic control. The result of this operation overall darkens the darker areas of the voxel-sets and lightens the lighter areas.

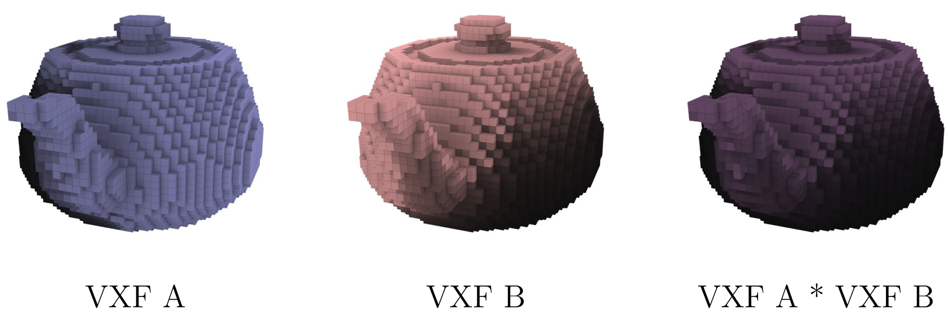

Figure 4.11: Blending Operations - Multiply

- Darker

  The darker operation selects the darker value between the illumination values of the primary voxel-set and the illumination values of the secondary voxel-set. An example of this operation can be seen in figure 4.12.

  The obvious result of this operation is a darker voxel-set which takes into consideration the dark areas of each of the two voxel-sets.

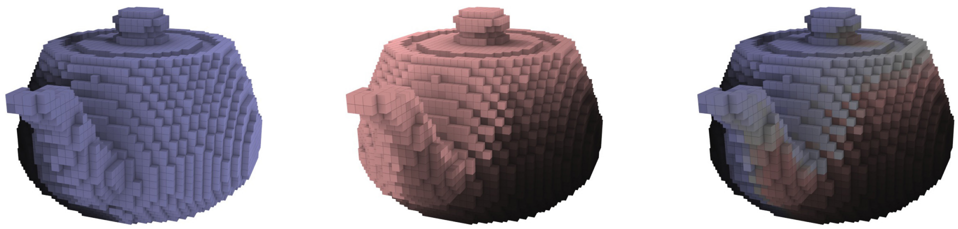

VXF A VXF B darker(VXF A, VXF B)

Figure 4.12: Blending Operations - Darker

- Lighter

  The lighter operation selects the lighter value between the illumination values of the primary voxel-set and the illumination values of the secondary voxel-set. An example of this operation can be seen in figure 4.13.

  The obvious result of this operation is a lighter voxel-set which takes into consideration the light areas of each of the two voxel-sets.

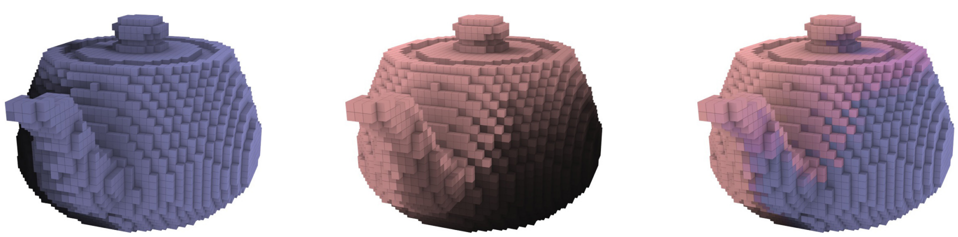

VXF A VXF B lighter(VXF A, VXF B)

Figure 4.13: Blending Operations -Lighter

4.2.2.2 Filter operations

The neighboring operations require a single voxel-set and perform different operations on each voxels based on its neighboring values. These operations can be automatic or can utilize a series of parameters external to the voxel-set.

Some of these operations tend to make use of the blending operations during their processing. These operations can be similar to common filters offered in common photo-manipulation packages.

- Blur

  This operation softens and blurs the illumination data in the voxel-set by averaging out illumination values in a certain neighborhood of a given voxel. An example of this filter operation can be seen in figure 4.14.

  The basic setting uses a box-filter to perform the blur operation, but also presents a flag that allows the influence of the voxels considered to be weighted by the distance from the voxel in consideration. The size of the neighboring area is a parameter assigned by the user when performing this operation. which averages the dark and light values and by result decreases the sharpness and the contrast present in the illumination values of the voxel-set.

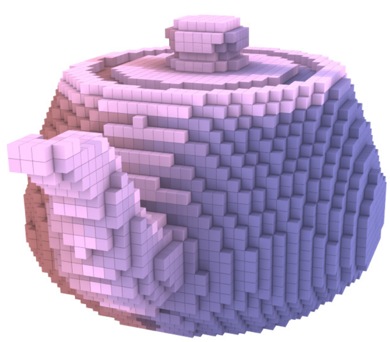

basic

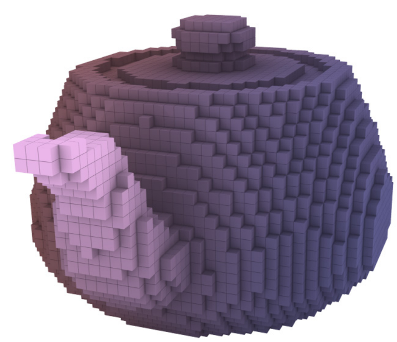

blur factor 10

Figure 4.14: Filter Operations - Blur

- Blur Special

  This operation was designed and implemented to achieve a very specific goal and look. The blur special operation is the key operation performed during the "Difference + smooth" export of the VXF interface tool.

This operation performs a form of blur to assign a newly computed illumination value only on active voxels with no illumination value or belonging to a pre-defined parametric group of voxels. An example of the blur special operation can be seen in figure 4.15.

This allows the operation to selectively not modify areas that have been defined by the user, while modifying other areas defined by the user. The blur special operation is used in the "Difference + smooth" export as it allows the user-defined value changes to bleed into the neighboring areas not touched by the user and to create a smoother transition between the original values and the user defined values.

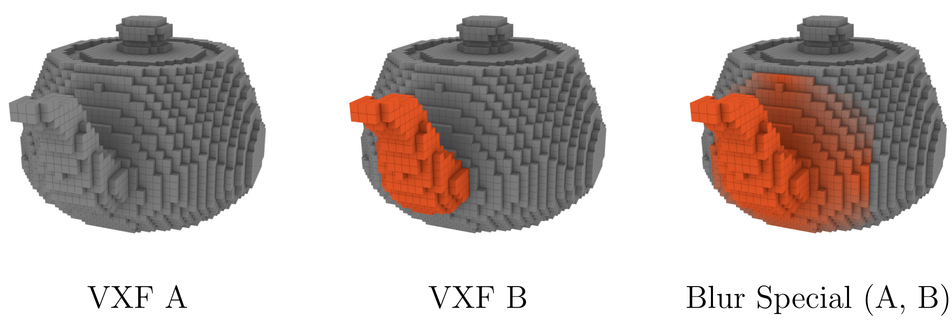

Figure 4.15: Filter Operations - Blur Special

4.2.2.3 Other operations

- Arithmetic operations.

This operation is very simple and its nature is very clear. The operation takes a single voxel-set and a single number and performs a simple arithmetic operation (addition, subtraction, multiplication, division and power degree) between each value of the voxel-set and the numerical term. This is not very useful with the most basic operations (addition and subtraction), but with the remaining operations (multiplication, division an power degree) it is quite useful as a helper operation in combination with blending modes and other more complex operations.

- Normalize.

  As the name clearly implies, this operation normalizes between 0 and 1 the illumination value of the voxel-set. This operation offers the possibility to use a flag that will set to 0 all negative values before proceeding to normalize the illumination values of the voxel-set. An example of this operation can be seen in figure 4.16.

  The normalized data set is useful for editing and artistic control purposes. The operation also declares the magnitude of the illumination data so it is possible to restore a normalized data set to its original values.

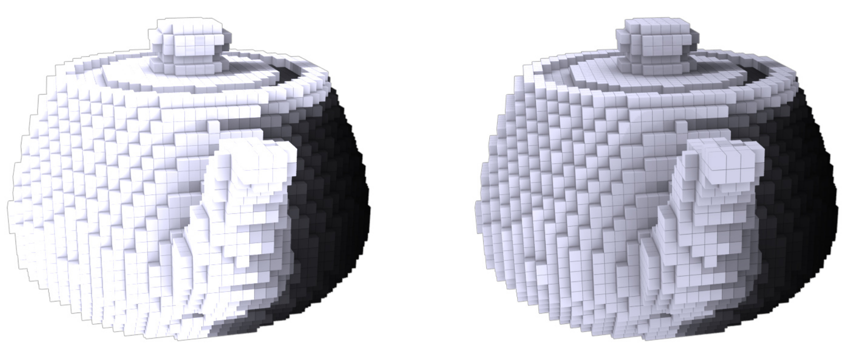

basic normalized

Figure 4.16: Other operations - Normalize

- Gradient operation. Given a voxel-set representing a volume, the gradient operations can generate a wide variety of voxel-sets based on basic gradients and directional ramps. These voxel-sets can be used in conjunction with other VXF operators and tools or independently. The gradient operator can generate a variety of voxel-sets:

  - Directional gradient. This type assigns value to the voxels based on two colors and a vector between two points.
  - Radial gradient This gradient assigns values to the voxels based on the distance from a single point, a falloff radius and two colors.
  - Distance from surface. This gradient assigns values to the voxels based on the distance of each voxel from the surface of the volume geometry.

  All gradients can also be extended to allow multi-value ramps and not exclusively ramps between two values.

- Remove / include. This is the only operation set that does not affect the illumination data stored in the VXF file. These operations take in two VXF files and voxel-sets and output a single voxel-set. The operations modify the active parameter of the single voxels.

  The "remove" operation will proceed to deactivate the voxels of the primary voxel-set that are indicated in the secondary voxel-set. This operation is performed exclusively when the voxel in the primary voxel-set is active.

  The "include" operation will activate all the voxels of the primary voxel-set that are not active and are indicated in the secondary voxel-set. This operation is performed exclusively when the voxels in the primary voxel-set are not active.

5. RESULTS

5.1 Overview

The overall result is an artistic toolbox which allows the user to generate, edit and manipulate the light transfer effects in participating media. In this following chapter some examples developed with this system are shown. The system allows the user to handle light and depth in a very dynamic and free manner.

While our method offers a linear workflow which evaluates the effects of light transfer, it also allows the artist extensive control over every single step of the process and even allows the artist to rearrange the steps freely and even forgo completely simulated evaluations in favor of other directives.

The analysis of the results is organized in two section. A first section which shows the artistic results achieved with our system, and a second section which will highlight the strength and weaknesses of our system.

5.2 Artistic results

5.2.1 Teapot

Figure 5.1: Teapot render

The teapot image, seen in figure 5.1, is the first example rendered with our new method. We have chosen a very traditional cg subject first our first test: a teapot.

Our method was used in a conservative way, attempting to recreate a relatively realistic effect and smoothness. The contribution of the scattering effect is conservatively balanced and not too extreme, as seen in figure 5.2.

The smoothness and the hues present in the scattering contribution pass already show some of the strengths of our method.

Scatter contribution — Direct contribution

Figure 5.2: Different passes that contribute to the final teapot image

5.2.2 Happybear

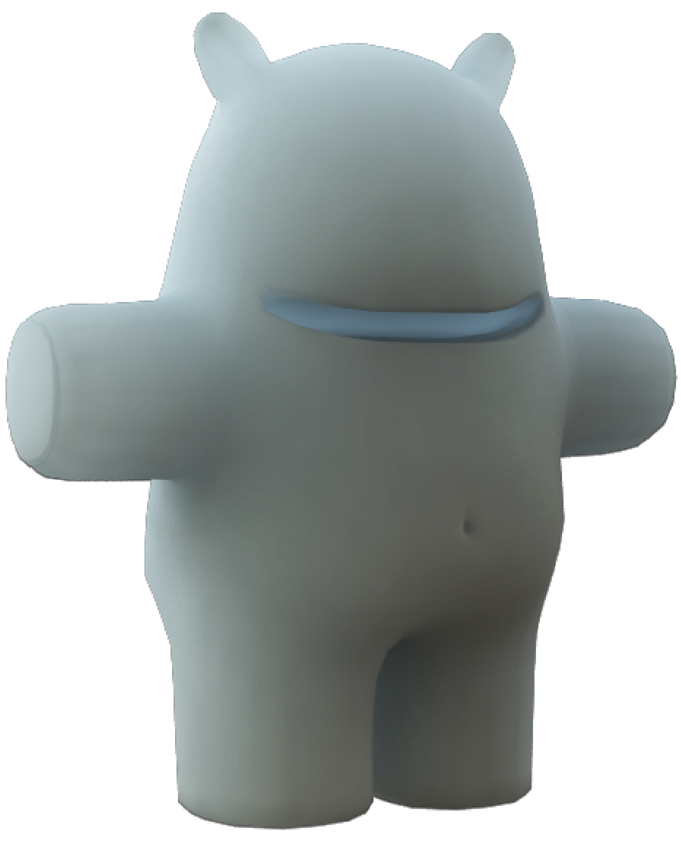

Figure 5.3: Happybear render

The renders of the "Happy bear" model show a slightly more extensive use of our method to achieve the desired visual effect, as seen in figure 5.3. The tools developed in this thesis were used to edit the illumination data gathered from the environment illumination. This allowed the recreation of different hues throughout the volume. We also used a gradient filter to increase the saturation and decrease the value of the colors inside the volume based on depth calculations. The different voxel-sets used to compose the final image can be seen in figure 5.4.

Environment scatter contribution

Gradient filter contribution

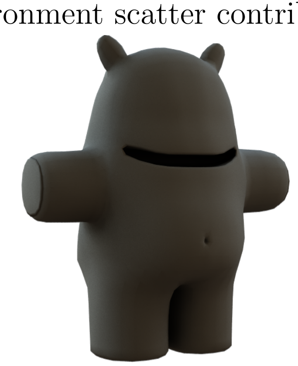

Direct contribution

Figure 5.4: Different passes that contribute to the final happy bear image

5.2.3 Rubber duck

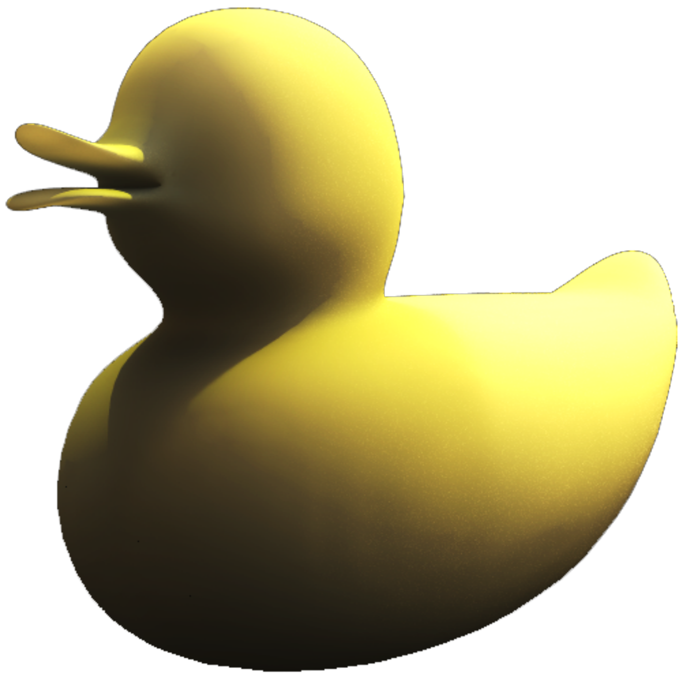

Figure 5.5: Rubber duck render

While aiming for a conservative visual effect, the rubber duck example, seen in figure 5.5, was created by using a multitude of VXF operators and tools. The scattering pass was edited to have smoother tones and also pick up different hues from the environment illumination. In this example the VXF tool was used to accentuate the illumination values of particular areas of the volume over others. Multiple different VXF sets were created and combined into a single scattering pass that is shown in figure 5.6

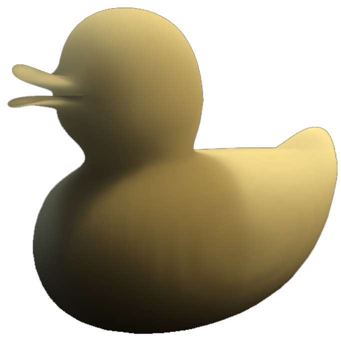

Artistically edited scatter contribution

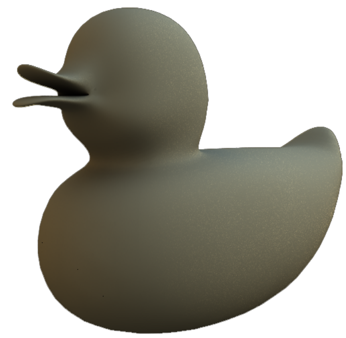

Environment contribution

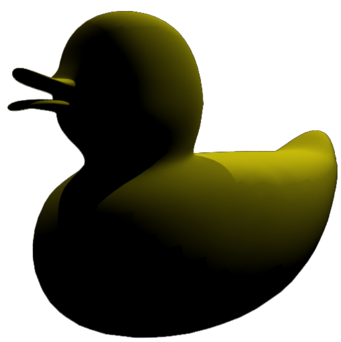

Direct contribution

Figure 5.6: Different passes that contribute to the final rubber duck image

5.2.4 Various artistic results

In figure 5.7, we show a variety of results that extensively use the toolbox offered by this thesis. These examples combine multiple VXF operations, a series of independently modeled VXF sets and multiple passes of layering values and intensities with the VXF tool, Though the tools might not be the most user-friendly, these images show the extent of the control of the visual effects of scattering. The toolbox offered with this method allows the artist to freely control hue, value and saturation of the effects.

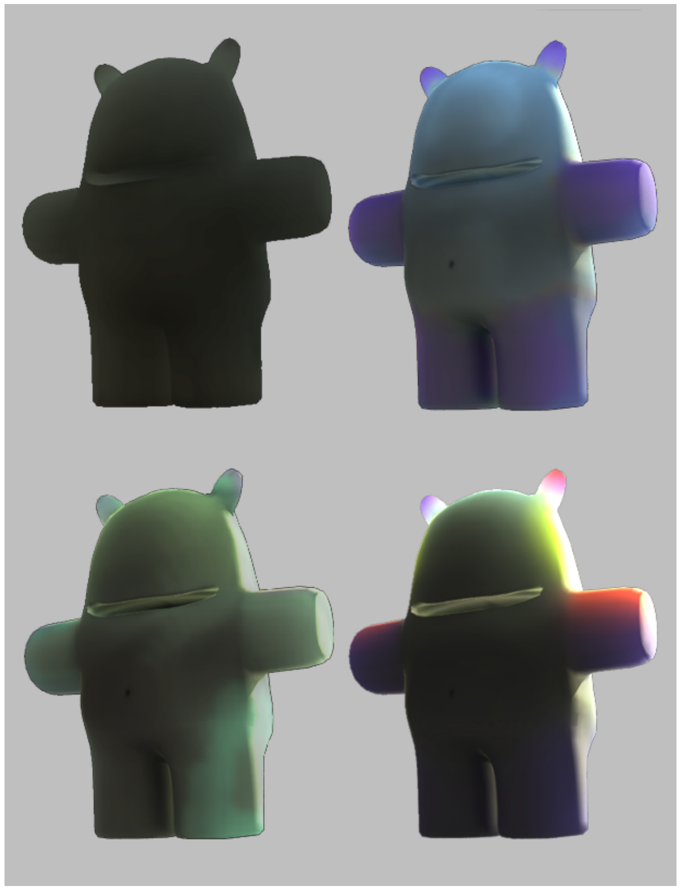

Figure 5.7: Various artistic interpretations using the various tools and our render method.

5.3 Technical limitations and issues

While the approach offers the artist several tools and allows the artist to recreate a wide variety of visual effects, after testing our approach a few technical limitation appear obvious.

Our method recreates the visual effects of scattering by approximating an isotropic phase-function. This is due to the assumption that the isotropic scattering is the dominant component of the visual effects of scatter. While this assumption allows the user wide control and variety of effects, there are several materials in which dominant scattering effects are not isotropic and therefore are difficult and impractical to recreate with our method.

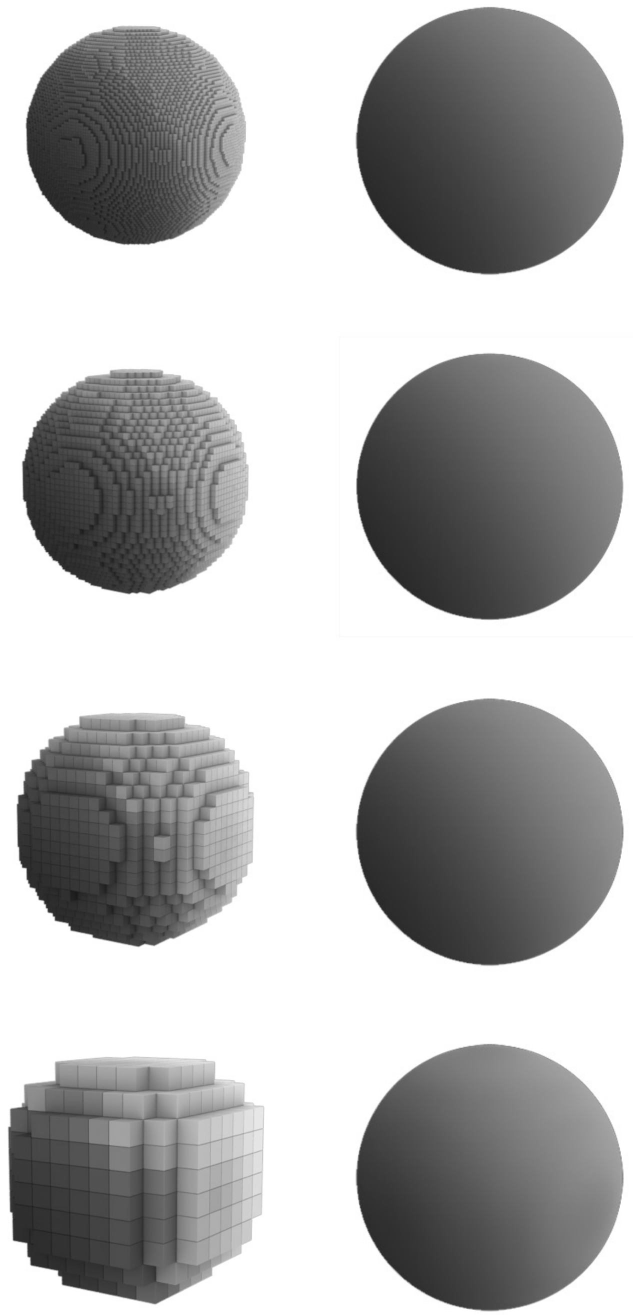

size 0.25

size 0.5

size 1

size 2

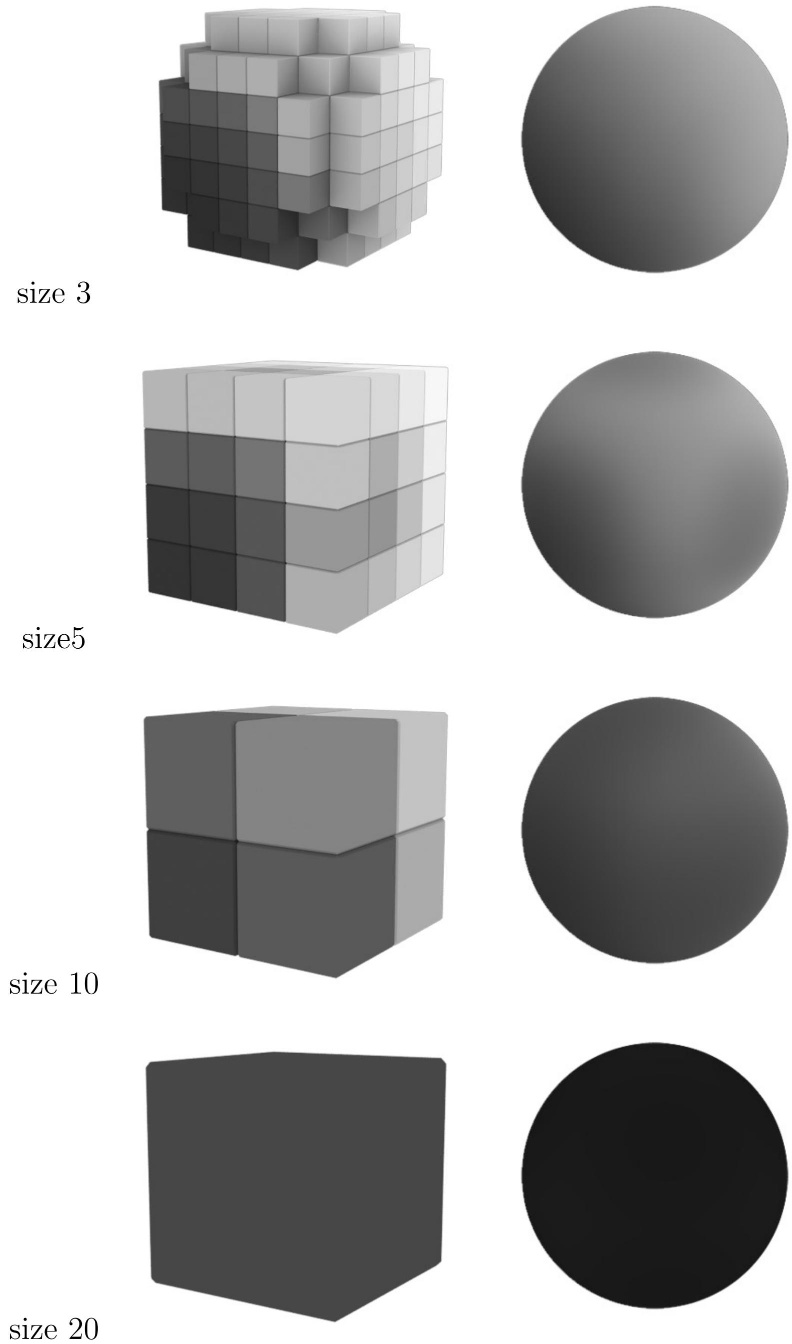

Figure 5.8: The effect of light transfer in participating media - cross section of teapot object

Our system presents two limitations tied to the technical implementation used.

The rendering method is based on a discrete representation of the volume and several operations acting on this discrete representation. Considering the central importance of this representation, it is fundamental that the voxel-set is able to represent the shape details of the surface in a consistent way. In figure 5.8, it is possible to notice how lower voxel resolutions can fail to recreate the visual results and cause some visual artifacts. These are simple examples with a very basic sphere. In figure 5.9 it is possible to observe a dark artifact at the tips of the ears. This area is darker due to the voxel resolution being too low. In situations where the voxel resolution cannot properly represent a variation of curvature or extreme curvature values, various visual artifacts can occur.

Figure 5.9: Scattering render with visible artifacts

These artifacts can be avoided by increasing the voxel resolution. However increasing the voxel resolution presents a huge computational hike and often reach exaggerated rendering time.

Another technical limitation inherent to our approach is a sampling issue. Our implementation traces through the volume to compute and gather the in-scattering contribution for each voxel. An inappropriate sampling resolution will not correctly approximate the scattering from neighboring voxels. Therefore it is necessary to select a correct sampling rate according to the voxel resolution.

6. CONCLUSION

The primary goal of this research was to develop a rendering method and a collection of tools and processes that would allow a user to artistically control and model the effects of light transfer in participating media.

This thesis is focused on developing an approach based on an approximated model that would allow the user to intuitively manipulate the illumination data and the results of light transfer.

In order to achieve this artist-friendly approach and develop a series of processes that find close analogues both in traditional artistic techniques and in contemporary digital artistic techniques, some physical attributes and behaviors were approximated. As expected, these approximations limit the inherent ability of our approach to automatically recreate the more minute and physically accurate details of light transfer in participating media; but they were necessary to achieve the artistic interface goals of the research.

The rendering method we developed successfully recreates the more recognizable visual results of light transfer, but the more minute visual details are left to the users and the set of tools offered by our approach.

The resulting approach is flexible and offers a wide selection of tools for easily editing and manipulating the illumination data and the final visual results of light transfer. It also offers a basic platform that allows for the expansion of the tool set.

The approach enables the user to create compelling images and carry out an artistic vision. While this is possible with the tools and processes offered by this thesis, it is fundamental to concede that the interface and these tools developed in this research are prototypes of ideas and concepts not fully developed and tested artistic tools. Nevertheless the results of our research clearly show potential for this approach and also suggest ideas to extend and improve from technical point of view, as well as artistic and user experience perspectives.

7. FUTURE WORK

The rendering method and the toolbox we have developed opens the door to various different projects. Our approach can be extended by tackling those technical limitations due to the experimental approach to the implementation. Developing a better organized and quicker data structure to represent our volume would increase the speed of iterations and favor artistic control. The current VXF data structure has no optimization and stores the volumetric data in an unorganized format which cause inefficient access to the data stored.

Another limitation intrinsic to the VXF data structure is the uniform resolution of the voxel-set. In the previous section we analyzed the artifacts and issues caused by inadequate voxel resolution. Increasing the voxel resolution helps avoid the visual artifacts, but increases the computational time and adds considerable amounts of redundant data to the VXF data structure. The option to have variable levels of detail and resolution in areas of the volume where needed and lower resolution where not needed would potentially eliminate the artifacts mentioned in the previous section, increase the speed during the rendering process and also increase the control of detail effects of scattering.

Taking into consideration the artistic toolbox section of this method, there are several opportunities to develop new tools and new interaction models to aid the user in his artistic goals. The current approach used to paint into the voxel-set is very mechanical and requires multiple setup steps to function correctly. The tool is burdened by some interface elements integral to Autodesk Maya and currently does not use Autodesk Maya's own painting tools to full extent. While the current tool does allow the artist to paint on any voxel at any depth or position in the voxel set, it is time consuming and has a very steep learning curve. The tool could also be extended to include many of the features already embedded in Autodesk Maya, potentially allowing the user to use 3d gradients, 3d noises and various surface and volumetric projection techniques available in that software package.

Other than the VXF tool running in Autodesk Maya, the current artist toolbox is composed by a wide selection of independent command-line tools. The great majority of these tools take a VXF file and perform some operations and output a new VXF file. During the testing and use of the toolbox, I had developed a series of one-off script to concatenate a series of operations. Due to nature of these tools, it should be relatively simple to develop a node-base tool that would allow the artist to create his own node graphs of multiple operations. The update of the VXF tool and the development of a node-based editor for VXF operators would boost the potential of the artistic tool set enormously and result in lesser technical struggles and more freedom to achieve the artistic vision.

The current system was developed with the focus on the visual effects of light transfer in solid volumes. Many of the same artistic and rendering principles could be used in other forms of volumetric rendering. Similar concepts could be used to control the visual effects of scattering through smoke, clouds or other types of volumes. The current rendering system also does not take into consideration the possibility that the viewpoint is inside of the volume. While not as relevant when considering solid volumes, if future developments of the system were to includefluids or gaseous volumes, this condition will have to be tackled.

Early on in the development, we decided to not rely on external renderers or applications to ensure freedom and control over each step of the system. With exclusion of the VXF tool, we were able to develop the entire system without relying on any application and hence allowed us to identify what our approach required. It would now be possible to develop our rendering method using an external rendering engine and still keep the degrees of control we have defined. Many of the modern renderers present various technical advantages and also offer a variety of volumetric rendering algorithms that could further enhance our method.

REFERENCES

[1] Ayn Rand. *The Romantic Manifesto: A Philosophy of Literature*. New American Library, New York City, NY, USA, 1969.

[2] Subrahmanyan Chandrasekhar. *Radiative Transfer*. Dover Publications, New York City, NY, USA, 1960.

[3] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 511–518, New York, NY, USA, 2001. ACM.

[4] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometrical considerations and nomenclature for reflectance. In Lawrence B. Wolff, Steven A. Shafer, and Glenn Healey, editors, *Radiometry*, pages 94–145. Jones and Bartlett Publishers, Inc., 1992.

[5] Frederic Pérez, Xavier Pueyo, and François X. Sillion. Global illumination techniques for the simulation of participating media. In *Proceedings of the Eurographics Workshop on Rendering Techniques '97*, pages 309–320, London, UK, 1997. Springer-Verlag.

[6] Eva Cerezo, Frederic Pérez, Xavier Pueyo, Francisco J. Seron, and François X. Sillion. A survey on participating media rendering techniques. *The Visual Computer*, 21:303–328, 2005. 10.1007/s00371-005-0287-1.

[7] Holly E. Rushmeier and Kenneth E. Torrance. The zonal method for calculating light intensities in the presence of a participating medium. *SIGGRAPH Comput. Graph.*, 21:293–302, August 1987.

[8] James T. Kajiya and Brian P Von Herzen. Ray tracing volume densities. *SIGGRAPH Comput. Graph.*, 18:165–174, January 1984.

[9] James T. Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '86, pages 143–150, New York, NY, USA, 1986. ACM.

[10] H.W. Jensen, J. Legakis, and J. Dorsey. Rendering of wet materials. In *10th Eurographics Workshop on Rendering*, pages 273–282. The Eurographics Association, 1999.

[11] Eric P. Lafortune and Yves D. Willems. Rendering participating media with bidirectional path tracing. In *In Eurographics Rendering Workshop*, pages 91–100. Springer-Verlag/Wien, 1996.

[12] Matt Pharr and Pat Hanrahan. Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '00, pages 75–84, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.

[13] Henrik Wann Jensen. Global illumination using photon maps. In *Proceedings of the eurographics workshop on Rendering techniques '96*, pages 21–30, London, UK, 1996. Springer-Verlag.

[14] Henrik Wann Jensen and Per H. Christensen. Efficient simulation of light transport in scences with participating media using photon maps. In *Proceedings*

of the 25th annual conference on Computer graphics and interactive techniques, SIGGRAPH '98, pages 311–320, New York, NY, USA, 1998. ACM.

[15] Julie Dorsey, Alan Edelman, Henrik Wann Jensen, Justin Legakis, and Hans Kohling Pedersen. Modeling and rendering of weathered stone. In *ACM SIGGRAPH 2006 Courses*, SIGGRAPH '06, New York, NY, USA, 2006. ACM.

[16] Simon Premože, Michael Ashikhmin, and Peter Shirley. Path integration for light transport in volumes. In *Proceedings of the 14th Eurographics workshop on Rendering*, EGRW '03, pages 52–63, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.

[17] S. Premože, M. Ashikhmin, R. Ramamoorthi, and S. Nayar. Practical rendering of multiple scattering effects in participating media. In *Proc. of Eurographics Symposium on Rendering*, 2004.

[18] Pat Hanrahan and Wolfgang Krueger. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '93, pages 165–174, New York, NY, USA, 1993. ACM.

[19] Jos Stam. Multiple scattering as a diffusion process. In *Eurographics Rendering Workshop*, pages 41–50, 1995.

[20] Craig Donner and Henrik Wann Jensen. Light diffusion in multi-layered translucent materials. *ACM Trans. Graph.*, 24:1032–1039, July 2005.

[21] Henrik Wann Jensen and Juan Buhler. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graph.*, 21:576–581, July 2002.

[22] Jeppe Revall Frisvad, Niels Jørgen Christensen, and Henrik Wann Jensen. Computing the scattering properties of participating media using lorenz-mie theory. In *ACM SIGGRAPH 2007 papers*, SIGGRAPH '07, New York, NY, USA, 2007. ACM.

[23] Wojciech Jarosz, Craig Donner, Matthias Zwicker, and Henrik Wann Jensen. Radiance caching for participating media. *ACM Trans. Graph*, 2008.

[24] Raanan Fattal. Participating media illumination using light propagation maps. *ACM Trans. Graph.*, 28:7:1–7:11, February 2009.

[25] Craig Donner, Jason Lawrence, Ravi Ramamoorthi, Toshiya Hachisuka, Henrik Wann Jensen, and Shree Nayar. An empirical bssrdf model. In *ACM SIGGRAPH 2009 papers*, SIGGRAPH '09, pages 30:1–30:10, New York, NY, USA, 2009. ACM.

[26] Wikipedia. Beer-lambert law — wikipedia, the free encyclopedia, 2011. [Online; accessed 11-January-2012].

[27] Wikipedia. Attenuation coefficient — wikipedia, the free encyclopedia, 2011. [Online; accessed 11-January-2012].