CHALMERS

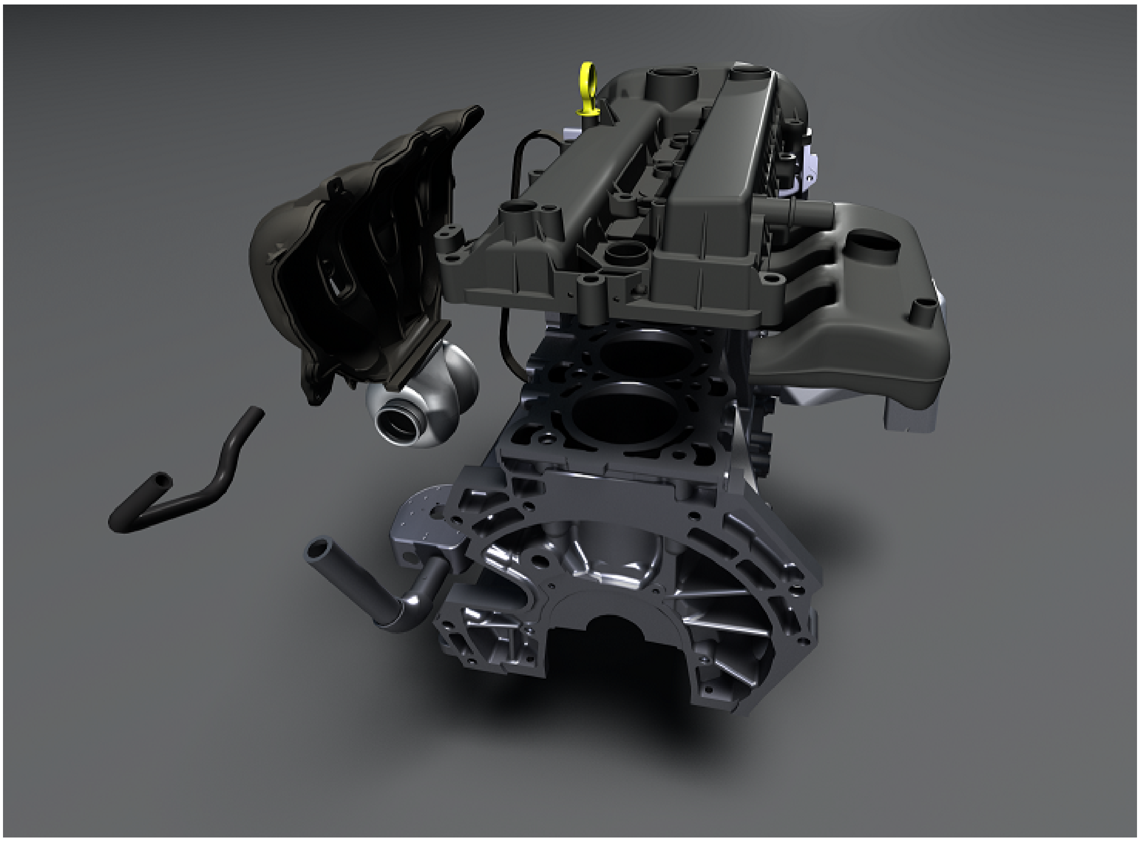

A practical shading model for ray tracing

Master's Thesis in Computer Science and Engineering

AKI KÄKELÄ

Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY AND UNIVERSITY OF GOTHENBURG
Gothenburg, Sweden 2016

A practical model for ray tracing
AKI KÄKELÄ

Supervisor: Erik Sintorn
Department of Computer Science and Engineering
Examiner: Ulf Assarsson
Department of Computer Science and Engineering

Computer Science and Engineering
Chalmers University of Technology and University of Gothenburg
SE-412 96 Gothenburg
Sweden
Telephone +46 (0)31-772 1000

Department of Computer Science and Engineering
Gothenburg, Sweden 2016

Cover: Example image produced with the implemented shading model and renderer.

Typeset in LaTeX

ABSTRACT

Photorealistic rendering relies on informative descriptions of its scenes. In particular, the material properties of 3D objects is important to achieve convincing results. The material of an object details how it interacts with light, for instance the color, transparency, and the sharpness of its reflections. What arises in practice is finding a brief description that allows for a wide range of materials to be represented, while preserving photorealism - in other words, a *practical* material system. These constraints mean that the material parameters should be few and easy to understand, and creating the material definitions themselves must be simple. This report outlines a shading model that attempts to deal with these issues, and presents a collection of base material types that exposes only a few, but powerful, parameters for artists to work with.

ACKNOWLEDGMENTS

I would like to thank both my supervisor Erik Sintorn and examiner Ulf Assarsson for providing helpful ideas and comments, and Vizendo and Olof Bjarnason in particular for continuous feedback and inspiring discussions.

Aki, Gothenburg, 8th October 2016

CONTENTS

LIST OF FIGURES

1 INTRODUCTION

1.1 Background

There are two principal methods of rendering computer graphics: *rasterization* and *ray tracing*. While rasterization is fast, it is more difficult to accurately simulate effects such as reflection, refraction, and shadows. Ray tracing naturally models these types of light interaction, but is often too slow or too noisy for real-time contexts. Ray tracing is used almost universally for CGI in movies, car catalogs, architectural visualization, and other media where visual quality is more important than interactivity.

Rasterization has always been favored in real-time settings due to high demands on interactivity and the highly optimized hardware, historically designed with rasterization in mind. Lately, GPUs have become more general-purpose through APIs including Optix (NVIDIA) and OpenCL (Khronos Group). Many rendering systems offload tasks and eventually implement rendering techniques on the GPU for faster rendering.

Both render methods rely on *shading models* to render images. A shading model can be thought of as the set of material properties that define how light interacts with an object, e.g. colors, transparency, or roughness, and the way in which these parameters are used when rendering. They are sometimes called *illumination models* or *lighting models* to contrast them from interpolation methods (e.g. Gouraud shading, Phong shading); a renderer will use a shading model to define object appearances and typically use interpolation techniques on the 3D geometry data. The shading model describes how light interacts with objects in a scene, and is therefore central in achieving high visual quality.

1.2 Application

The goal is to improve the visual quality in a virtual training application at the company Vizendo. The domain is mainly the automotive industry, where users of the application learn the proper assembling operations of manufacturer products. Training scenarios consist of using the correct tool (e.g. a screwdriver) with the correct item (a set of screws) on the correct parts in a particular order (mounting the screws in a predetermined pattern). The sequence is always linear: there is only one correct path in the procedure. For each action there is a corresponding animation, in addition to visual highlighting of the current step. The application also contains an editor, where the sequences are constructed and scene setting is done, including the mapping of materials to plain 3D geometry.

A rasterizer is currently used, which renders in real time to interactive rates, depending on the complexity of the scene. Ray tracing is well-suited since the sequences are predetermined and can therefore be prerendered. Another motivation is that the prerendered results are more easily played back on hardware less capable for 3D-rendering, such as tablets and phones. Potentially sensitive, proprietary manufacturer data (mainly geometry) is also stripped in the process, allowing for simpler distribution. The disadvantage of prerendering is the lack of free camera movement, which is typically not as important during training scenarios, as the camera locations are part of the predetermined training sequences.

1.3 Purpose

This work consists of creating a shading model for use in ray tracing that satisfies the needs of rendering with high visual quality within the domain. The renderer implementing this model must be able to handle it appropriately. A prototype ray tracer has therefore been implemented to demonstrate the shading model.

1.4 Problem

A more sophisticated rendering model is able to make use of more detailed data about the appearance of objects in a scene. Unfortunately, when it comes to photorealistic ray tracing, there is no clear standard for how this data should be represented. This can be explained by the potentially increased complexity of the rendering model compared to rasterization. For real-time rendering there are relatively simple, but standard, shading models that are frequently used, typically based on the Blinn-Phong shading model. As an example, a common format for defining 3D scenes that is influenced by Blinn-Phong is the Wavefront format for geometry (OBJ) and materials (MTL - Material Template Library). Blinn-Phong is a simplified model used for its low computation cost, and not particularly suitable for photorealistic rendering.

The aim is to develop a shading model that is practical to use by artists, while being suitable for ray tracing and closer to photorealism than the standard real-time models. A set of material types should be defined. The model should then be implemented in a prototype ray tracer. A *practical* shading model implies materials that are easy and quick to define. This sets a limit on both the number of parameters and their complexity.

1.5 Limitations

Although not the purpose of this work, a subgoal has been to explore the possibility of using ray tracing in an interactive setting, such that it can be used in the editor also. This would make scene setting easier due to faster iteration times. Performance is otherwise not the topic of this thesis work.

Both the renderer and the shading model focus specifically on the domain, which in practice is the car industry, and typical materials that occur in that domain.

The current system in place for geometry and materials rarely utilizes some aspects of rendering, for instance textures and texture mapping. Textures are an additional workload and are rarely provided along with the 3D data. The lack of textures means that the rendered object surfaces will have a spatially uniform macroscale appearance.

2 THEORY

This section will give an overview of the theory behind ray tracing and surface shading.

2.1 Ray tracing

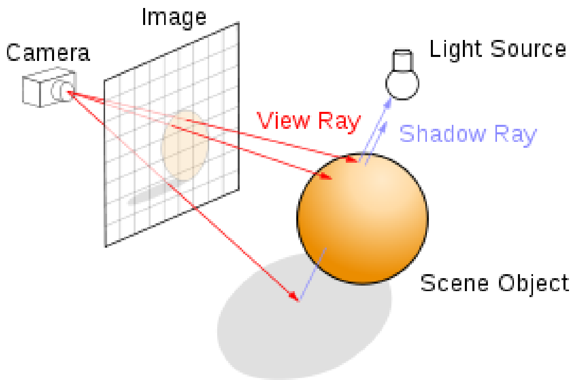

Figure 1: An example of tracing rays from a camera, through an image grid, into the scene.

Ray tracing involves sending rays from the camera, or eye, in the direction of the pixels on the screen (see Figure 1). These rays may reflect on or transmit into objects, or become absorbed. Colors are computed by tracing a ray for each pixel. The initial rays, originating at the camera, are called *primary rays*. In general recursive ray tracing, or *Whitted ray tracing* [1], up to three types of *secondary rays* are spawned at each intersection: reflection, refraction, and shadow rays. Shadow rays are cast toward each light source to determine whether the surface point is occluded toward that light source. If not, the light is sampled for its contribution.

There are several variants of ray tracing, with unique strengths and weaknesses. What they have in common is that they are more physically accurate than the raster model, but also require more computation. Shadows, for instance, will always be pixel-perfect, limited only by the detail of the geometry. The secondary rays compute indirect illumination, or *global illumination*, where for instance a red sphere bleeds its color onto a white floor, even though neither of them is a light source. How the light bounces and what colors are involved in the computations depend on the material properties of the hit surfaces.

2.1.1 *Path tracing*

First introduced by Kajiya [2] along with the light transport equation (see Section 2.2 on the next page), path tracing is an unbiased Monte Carlo algorithm for ray tracing. Instead of tracing rays in a tree, this technique builds single paths incrementally as the rays intersect with objects in the scene. Paths may have multiple stopping criteria, such as losing most of their energy, or that no geometry at all is intersected.

One sample consists of only one path per pixel, so as to avoid computing the less significant rays deep in the tree of general ray tracing. For photorealistic rendering without much noise it may be required to trace thousands of samples to reach an acceptably low level of noise in the image. Under-sampling results in grainy images due to the random nature of light scattering.

2.2 The light transport equation

The light transport equation (LTE; see Equation 1), also known as the *rendering equation*, gives the ratio of reflected radiance from direction w_r and incident irradiance from direction w_i at a point x. It is expressed as the sum of self-emission L_e, which is often zero, and the reflected radiance over all incident directions w_i of the hemisphere Ω above the point x. L_i is the incident illumination, modulated with the cosine law ($\cos\theta_i$), and f is the BSDF, described in Section 2.2.1. L stands for radiance, which is defined as incident power, or *flux*, per solid angle per projected unit surface area ($\frac{W}{m^2 sr^2}$).

This equation is difficult to solve analytically, [3] so numerical methods are employed instead. The path tracing algorithm is one such example.

$$L_r(x, \omega_r) = L_e(x, \omega_r) + \int_\Omega f_r(x, \omega_i, \omega_r) L_i(x, \omega_i) \cos\theta_i \, d\omega_i \tag{1}$$

2.2.1 *BSDF*

The LTE contains a term *f* which is the *bi-directional scattering distribution function* (BSDF). It describes the ratio of incident illumination from direction ω_i to reflected (or transmitted) illumination along direction ω_r at a point x. *Scattering* models (BSDF) encompass both reflection (B*R*DF) and transmission (B*T*DF). The term BxDF is used to refer to a function f with any of these types of interaction.

2.2.2 *Physically based reflectance models*

If the BSDF has particular properties we call it physically based: it must be (i) reciprocal ($f(x, w_i, w_o) = f(x, w_o, w_i)$), meaning that the incoming and outgoing directions can be swapped without affecting the value of *f*, and it must be (ii) energy-conserving, meaning the integral must evaluate to one or less. This is because outgoing energy cannot exceed incoming energy, but some energy may be absorbed. The value of *f* must also be (iii) non-negative.

2.3 Components of a reflectance model

Many light scattering models exist in computer graphics. The differences boil down to ease of use, evaluation cost, and correctness. This section will deal with components of typical reflectance models, with ray tracing in mind.

2.3.1 *Diffuse and specular reflectance*

Two extrema of *ideal*, mathematically simple models are diffuse and specular reflections. Ideally diffuse reflectance means that the reflectance distribution is uniform throughout the hemisphere above the surface. This type of reflection is sometimes called *body reflectance*, and it has a significant visual impact, often determining the majority color of a surface. It does so by approximating more complicated effects. One such effect is subsurface scattering, where light scatters underneath the surface and the point of exit may not be the same as the point of entry. Subsurface scattering has a smoothing effect, and may even yield different colors if there are sub-layers to the material.

As subsurface scattering is often prohibitively expensive to compute, body reflectance is normally approximated with a *Lambertian* BRDF, which is another name for ideally diffuse reflectance. In practice, few materials are ideally diffuse - the reflectance from a point often depends on at least the reflected direction, and the angle of incidence. Section 2.3.5 on page 8 describes a common model for this purpose.

The distribution function for ideally *specular* reflectance is a Dirac delta function: for all values other than the ideal reflected direction, the function evaluates to zero, with an integral of one over the hemisphere. In other words, the distribution has a singularity at the ideal reflected direction, which is simply the incoming direction mirrored about the surface normal[1].

Surfaces can also be glossy, which is anything between diffuse and specular. In addition, they may have biases toward other directions than the specular direction. One such example is retro-reflection, where light tends to scatter back along the incoming direction instead. This is a property of the surface of the moon, and various fabrics such as velvet and satin.

2.3.2 *Transmittance*

Transmission is the effect of light penetrating a surface, typically bending according to Snell's law as illustrated in Figure 2 on the next page. The effect depends on the index of refraction for the two media, for instance air and glass. Transmission may be defined either in terms of surface interaction only (allowing a ratio of light to pass through a boundary) or as decaying transmission defined by the absorbance and thickness of

1 A surface normal is a vector perpendicular to the tangent plane of the surface, pointing out of its front face.

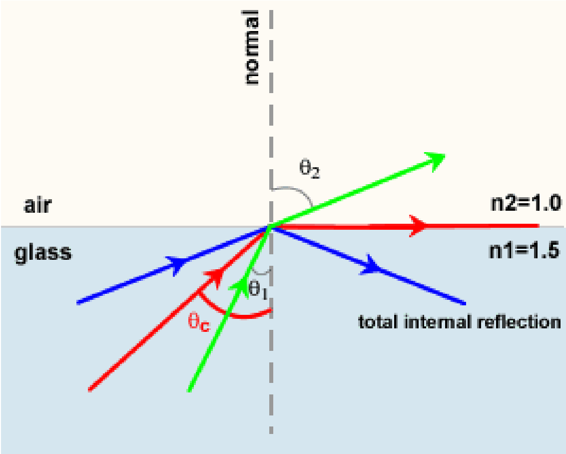

Figure 2: Reflection and refraction. Total internal reflection may occur when light travels from an optically denser medium.

the medium. This distinction becomes importance when comparing objects of different thickness: according to the Beer-Lambert law, thick media absorb more light. A surface-only transmittance factor is cheaper to evaluate, and may be simpler to use, but is much less accurate.

2.3.3 *Index of refraction*

The index of refraction is a complex number where the real part signifies the speed of light through the medium and the imaginary part the absorption (or extinction) of light by the medium. For non-metals, the imaginary part is often ignored, and the real part is further approximated with a single, average value over the entire spectrum of visible light, e.g. a value of 1.5 for glass. Using wavelength-dependent indices, for instance by having a refractive index for each of red, green, and blue, allows for rendering subtle color shifts, which is most significant in metals.

The refractive index also plays a crucial role in the Fresnel equations for computing the reflected and transmitted portions of light (described in Section 2.4.2 on page 10).

2.3.4 *Surface roughness*

Surface roughness has the effect of causing light to scatter in more varied directions. Optical surface roughness significantly affects the appearance of objects. All surfaces have roughness features to some degree, anywhere from the atomic scale and up. When these irregularities are smaller than the wavelength of visible light the surface will appear smooth, although diffraction still occurs. Diffraction has the effect of faintly warping the light's path around obstacles. Wave properties of light are usually ignored in rendering, as their effects are generally not noticeable. Only the optical laws of light are applied.

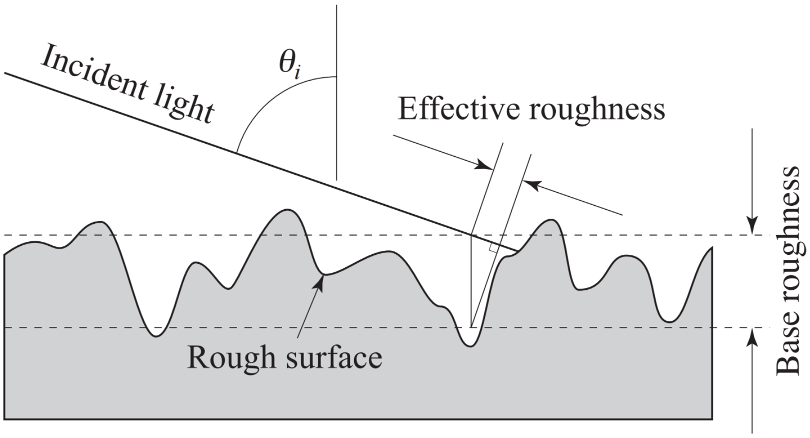

Figure 3: A grazing angle can make a rough surface appear smooth. (Image from [4].)

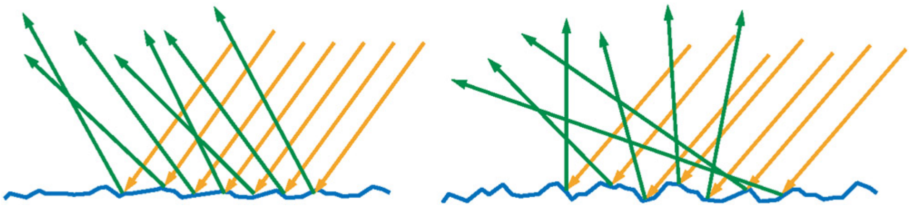

Figure 4: Surface roughness modeled with microfacets. The series of facets, shown in blue, are microscale and modeled statistically with distribution functions - not by explicit geometry. (Image from [5].)

Even rough surfaces will appear smooth at grazing angles, as shown in Figure 3.

2.3.5 *Microfacets*

Microfacet theory attempts to model rough surfaces by describing them as a series of V-shaped, microscopic cavities, where each *facet* has a particular slope (see Figure 4). In other words, each facet will have a particular surface normal. In the context of the macro scale, meaning detail at least on the level of a pixel on the screen, these microfacets are viewed as a whole by a statistical distribution of normals. This avoids both the modeling of actual micro geometry, and the computational expense that goes with it.

The *Normal Distribution Function* (NDF) describes the ratio of facets that have their normals perfectly aligned with a particular vector - this is the only case where a microfacet will reflect light. In other words, facets are *ideally specular*. In addition to reflecting light, facets can interact by either (i) occluding incoming light (self-shadowing) or by (ii) blocking light reflected off other facets (masking). These effects diminish the light either between the incoming direction and the surface, or the surface and the outgoing direction. In theory, microfacets could also reflect light between themselves, but this is normally ignored - instead we use *single-scattering* models. Multiple-scattering microfacet models exist, [6] though the computational expense is significant.

2.4 Microfacet reflectance models

Microfacet models in computer graphics stem from the Torrance-Sparrow model [7], which was further developed with the Cook-Torrance model [8]. More recently, alterations to Cook-Torrance have been explored and have become popular. This section will review the latter two, beginning with a description of common terms.

In general, to describe the microfacet model f in terms of reflection only (that is, the BRDF), we have:

$$f = k \times f_{spec} + d \times f_{diff}, \tag{2}$$

where f is the BRDF result, f_{spec} is the specular portion of light, and f_{diff} the diffuse. For energy conservation we also have:

$$k + d = 1. \tag{3}$$

The specular term is defined as:

$$f_{spec} = \frac{F \times G \times D}{norm} \tag{4}$$

There are various configurations for the choice of F, G, and D terms, and sometimes also in the normalization factor *norm*. The models typically describe only the specular component f_{spec}. The diffuse component f_{diff} may for instance be set to ideally diffuse (Lambertian), or even set to zero for some materials.

2.4.1 *Definitions*

Since these models use microfacets, we use the microfacet normal m rather than the geometry normal n for much of the shading.

h is the half-angle (or half-way) vector defined as:

$$h = normalize(w_i + w_o), \tag{5}$$

thus it is half-way between the incident direction w_i and the scattered direction w_o. The NDF will be evaluated in terms of h, meaning only facets aligned perfectly along h will contribute.

r is a parameter to the NDF and represents the roughness of the surface. A higher value of r gives the surface a duller appearance, as the distribution will yield a smaller portion of facets aligned with the half-way vector. A roughness value of zero corresponds to a perfectly polished surface. Microfacet models are only valid for rough surfaces, so this case reverts to the simpler specular reflection (or transmission) model instead.

2.4.2 *Fresnel term F*

The Fresnel term describes the reflectance of a surface, given the indices of refraction for the two media. In practice, one of the media is often air, with light traveling either from or into another medium. Due to energy conservation we can assume that light which is not reflected is instead transmitted, if absorption is negligible: $1 = T + R + A$ with $A = 0$ yields $T = 1 - R$.

The full Fresnel equations are defined here without taking into account polarization of light. We assume all light is either perpendicular R_s (from the German word *senkrecht*) or parallel R_p. There is a simplified version for dielectrics that is less computationally expensive and uses only the real part of the refractive index, and simplifies it even further by averaging it to a single value over all wavelengths. For metals, color shift is much more apparent, and the Fresnel equations may be computed for each of N wavelengths, e.g. $N = 3$ for red, green, and blue.

In either case, we have the total reflected light R defined as the average between perpendicular-polarized and parallel-polarized light, assuming equal quantities of both such that light is unpolarized:

$$R = \frac{R_s + R_p}{2} \tag{6}$$

An approximation for the Fresnel equation for conductors (metals) is [3]:

$$R_p = \frac{(\eta^2 + k^2)\cos\theta_i^2 - 2\eta\cos\theta_i + 1}{(\eta^2 + k^2)\cos\theta_i^2 + 2\eta\cos\theta_i + 1} \tag{7}$$

$$R_s = \frac{(\eta^2 + k^2) - 2\eta\cos\theta_i + \cos\theta_i^2}{(\eta^2 + k^2) + 2\eta\cos\theta_i + \cos\theta_i^2} \tag{8}$$

And simplified for non-metals:

$$R_p = \frac{\eta_t \cos\theta_i - \eta_i \cos\theta_t}{\eta_t \cos\theta_i + \eta_i \cos\theta_t} \tag{9}$$

$$R_s = \frac{\eta_i \cos\theta_i - \eta_t \cos\theta_t}{\eta_i \cos\theta_i + \eta_t \cos\theta_t} \tag{10}$$

In the equations, i stands for the incident medium, and t for the transmitted medium - the medium that the ray travels into. η is the index of refraction with k being the imaginary part, and θ is the angle of incidence from the surface normal.

In real-time rendering, the Fresnel equation for dielectrics is sometimes further approximated as [9]:

$$F = f_0 + (1 - f_0)(1 - \cos)^5\theta \tag{11}$$

where:

$$f_0 = \frac{\eta_1 - \eta_2}{\eta_1 + \eta_2} \tag{12}$$

We opt to use the full equations instead of this approximation. All microfacet models presented here use the same F term.

2.4.3 *Geometry term G*

The geometric attenuation term accounts for the shadowing and masking effects of microfacets, described in Section 2.3.5 on page 8. The two variants, Blinn and Smith, will be described in the sections to follow.

2.4.4 *Distribution term D*

The distribution term (NDF) gives a statistical distribution of the orientations of the microfacets. Given a vector, it returns the proportion of microfacets whose normals are aligned precisely along that vector. Importantly, microfacets only contribute reflected light when aligned this way, since microfacets are ideally specular.

The distribution term often has the largest impact on the end result. It dictates the general shape of the specular highlights, in both size and fall-off.

Additionally, distributions can be biased so as to represent skewed reflections as found in brushed metals or CDs. The shape of the highlights is therefore also dependent on the NDF. In more precise terms, the roughness is defined by two values representing roughness along the vectors of the tangential plane defined by the normal. A rotation parameter in the tangential plane allows for arbitrary skewing. These anisotropic versions will not be presented.

2.4.5 *Cook-Torrance*

The Cook-Torrance microfacet model is implemented as follows:

$$\begin{aligned} G_{Blinn} &= \min\{1, \frac{2(n \cdot h)(n \cdot w_i)}{(w_i \cdot h)}, \frac{2(n \cdot h)(n \cdot w_o)}{(w_i \cdot h)}\} \\ D_{Beckmann} &= \frac{e^{-(\frac{\tan\theta_m}{r})^2}}{r^2 \cos^4\theta_m} \\ norm &= \pi(h \cdot w_i)(h \cdot w_o), \end{aligned} \tag{13}$$

where θ_m is the angle between the micro surface normal m and the half-vector h. (We can represent $\cos\theta_m$ as the dot product $m \cdot h$.)

The Cook-Torrance model can apply various other distribution terms. The term shown in Equation 13 is the Beckmann distribution, which has long been standard for microfacet models. An alternative pointed out in their paper from 1982 is the Gaussian distribution.

2.4.6 *GGX*

The Cook-Torrance model contains separate terms to deal with different phenomena of light interaction. In particular, the D and G terms can well be replaced. Walter et al. [10] introduced a new D term, which they dubbed GGX, and chose a different geometry term G (Smith). The normalization factor is also changed. The paper implements the evaluation as follows:

$$
\begin{aligned}
G(w_i, w_o, m) &= G_1(w_i, m) G_1(w_o, m) \\
G_1(w, m) &= \chi^+\left(\frac{w_i \cdot m}{w_i \cdot n}\right) \frac{2}{1 + \mathrm{erf}(a) + \frac{1}{a\sqrt{\pi}} e^{-a^2}} \\
D_{Beckmann} &= \chi^+(m \cdot n) \frac{e^{-\left(\frac{\tan\theta_m}{r}\right)^2}}{\pi r^2 \cos^4 \theta_m} \\
D_{GGX} &= \frac{r^2 \chi^+(m \cdot n)}{\pi \cos^4 \theta_m (r^2 + \tan^2 \theta_m)^2} \\
norm &= 4(h \cdot w_i)(h \cdot w_o)
\end{aligned}
\tag{14}
$$

where:

$$
\chi^+(a) = 1 \text{ if } a \geqslant 0, \text{ and } 0 \text{ otherwise.} \tag{15}
$$

Note also that the Beckmann D has a different form of evaluation; the discrepancy is not commented on in the paper. The principal change is in introducing the new distribution, GGX, which is derived from the Smith G [11] term, which in turn is an approximation to shadow-masking. In other words, the G and D terms are related. These changes lead to longer tails on specular highlights, and a closer fit to real-life measured data, at least of various glass surfaces as presented in the paper. A performance note is that the erf in the G_1 function is expensive to compute - instead, a rational approximation (not shown here) is often used, that has a low relative error ($< 0.35\%$).

GGX uses a normalization factor of 4 instead of π, as it supposedly leads to a closer fit to measured data.

The paper also describes the full microfacet model for transmission. It is put to use in the shading model, but its formulation will be omitted here.

2.5 Material classes

Materials can be classified well by the way they interact with light. A natural separation is whether the material is a metal (conductor) or a non-metal (dielectric). These two types of materials have considerably different reflectance properties. For instance, metals do not exhibit subsurface scattering - all light is either absorbed immediately, or reflected. At normal incidence - that is, looking head on at the surface - metals reflect

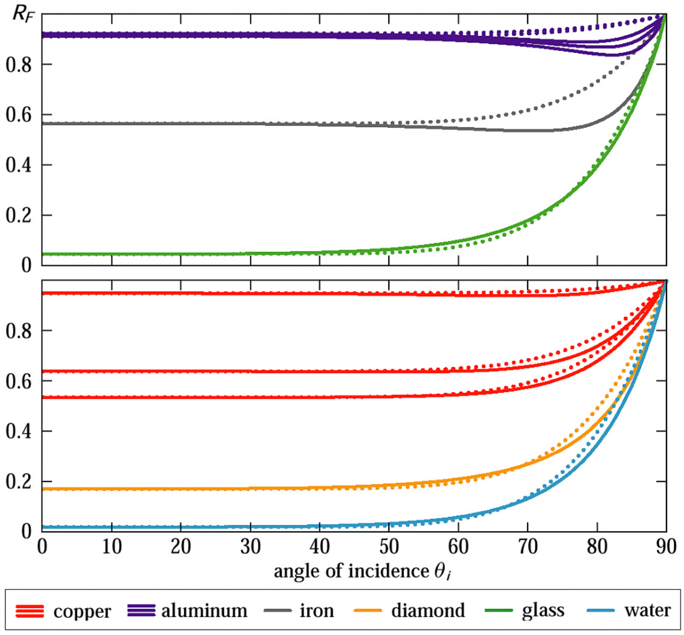

Figure 5: Reflectance as a function of angle of incidence. When looking head-on at a surface ($\theta_i = 0$), metals reflect significantly more than non-metals. Diamond is an outlier and one of the most reflective dielectrics outside laboratory environments. (Image from [5].)

significantly more light than non-metals, as illustrated in Figure 5. At grazing incidence, both classes reflect nearly all light, even though the surface may be rough.

Air, ice, glass, gemstones, rubber, and plastics are examples of dielectrics.

3 PROTOTYPE RENDERER

This section will discuss the choice of integration method for solving the integral in the LTE and give an overview of the capabilities of the prototype renderer.

3.1 Requirements

The typical scenes occurring in the domain set the requirements for what the renderer must be able to handle, but also give some leeway. For example, scenes will always consist of a centered object lit by surrounding lights. Some effects, like caustics, are also less important since glass material is typically only found on the exteriors of objects, e.g. the window or headlights on a car.

Ideally, rendering results should be displayable as soon as possible for shorter iteration times. Certain computationally heavy effects also have correspondingly low priority. In particular, it is desirable to get a good idea of the final image.

These criteria suit the path tracing algorithm well. Other techniques may produce blotchy specular reflections (photon mapping), or may suddenly introduce distinct detail into the image for caustics (Metropolis Light Transport). Path tracing also does not require any preprocessing, and the results can be displayed after each attained image sample. Its main weaknesses are scenes with significant indirect illumination and converging on caustics, and small light sources in particular. Caustics is considered less important, and typical scenes are directly lit, making the direct lighting weakness less meaningful.

3.2 Intersection tests

The renderer internally uses Intel's Embree ray tracing core [12], which is a highly optimized library for both ray queries on plain 3D geometry and for constructing the speed-up structures themselves. The library provides both closest-intersection tests and general occlusion tests, for determining either the nearest occluding object or whether any object at all is intersected. The latter is useful for shadow rays. Intersection queries provide information about local hit coordinates and the ID of the geometry that was hit. The library is otherwise unrelated to shading and ray tracing algorithms.

Embree also allows for user-defined geometry, given that bounding, intersection, and occlusion functions are provided. In this manner, other types of primitives can be supported in intersection queries. The renderer makes use of spheres and triangle meshes.

3.3 Input

There is a proprietary format for loading scene geometry and materials, which are imported from a scene graph structure. The renderer also reads the OBJ format.

3.4 Lights

As in photography, lighting in a renderer is important. The lights themselves do not affect the materials, but assist in demonstrating their effects. Different types of light sources are used to demonstrate the shading model: point lights, spherical lights with and without associated geometry, and uniform environmental lights. The uniform environmental light has constant radiance and surrounds the scene.

3.5 Direct illumination

Direct and indirect illumination are computed separately. The color value for each ray is accumulated by evaluating the contribution of directly visible light sources. If a geometric light is hit, its contribution is added directly.

If no geometry is intersected, the color value is found either from a background texture if it was a primary ray, or by querying the environment lights, if any.

3.6 Indirect illumination

Rays that intersect will often reflect on the surface of the object. This reflected direction depends on the material properties of the surface. In path tracing it will be a single direction, so good sampling is crucial to reduce variance. For example, an ideal mirror will always produce the ideal reflection direction, while an ideally diffuse surface will sample the entire hemisphere uniformly. For more involved materials, such as the BRDFs presented in Section 2.4 on page 9, the outgoing direction is importance sampled on the BSDF term. A microfacet model cannot in general be sampled exactly [10].

4 SHADING MODEL

The domain limits the space of materials that should be representable with the shading model. The most important materials are metals and metallic paint, plastics and rubber, and glass. Interiors also often contain various fibrous materials, though macroscopic spatially varying features like these are more suitably represented with textures, which the editor makes little use of. Textures will therefore not be a part of the model - this includes normal maps and displacement maps, but also spatially varying material properties such as base color or surface roughness, or textured metallic flakes in car paint.

Materials are defined in terms of reflectance models: any combination of BRDFs, BTDFs, and BSDFs. Since this can quickly get out of hand, and since energy conservation can become difficult to preserve, a select group of these combinations are presented outward. These are matte, mirror, glass, metal, plastic, and car paint. This section will describe both the BxDFs and the materials defined by them.

4.1 BxDFs

BxDFs are the building blocks of materials. Each one has to deal with energy conservation, such that compositions of these models will also conserve energy. This section presents the BxDFs of the model.

4.1.1 *Lambertian*

The Lambertian reflection model is perhaps the simplest one - it models perfectly diffuse reflection. It consists of a single parameter, which is the reflectance factor, or color. Variance occurs due to each direction in the hemisphere contributing equally.

4.1.2 *Oren-Nayar*

Oren-Nayar [13] developed a microfacet model where, instead of facets being ideally specular, they are ideally diffuse. It generalizes the Lambertian model to account for certain materials where roughness is inherent as in clay, plaster, and sand. It has the effect of dulling the highlights toward grazing angles.

The simplified version is used, also called the qualitative model. It ignores the small contribution of interreflections among facets. Oren-Nayar simplifies to the ideally diffuse Lambertian model when the roughness parameter is zero.

4.1.3 *Ideal specular reflection and transmission*

Both reflection and transmission are simple to calculate in the ideal case where we do not consider for instance the refractive index of the surface. In reflection, we always choose the ideal reflection direction $w_r = 2(n \cdot w_i)n - w_i$. For transmission we simply let the ray pass through without changing its direction: $w_t = -w_i$. No energy is lost in this simplified case.

The probabilities of choosing these directions are always one: there is a singularity in the distribution. This also means that there is no variance when sampling.

4.1.4 *Dielectric specular reflection and transmission*

The difference to the ideal case is that the Fresnel equations for dielectrics are used (see Section 2.4.2 on page 10), hence there is a ratio of reflected and transmitted light for a given angle of incidence. The ray is also refracted during transmission, due to the relative index of refraction between the two media on either side of the surface of the object. During transmission, total internal reflection may occur, as illustrated in Figure 2 on page 7.

4.1.5 *Conductor specular reflection*

As conductors do not transmit light, only a specular reflection BxDF is provided. This is similar to the dielectric case but uses the wavelength-dependent Fresnel equations instead.

As noted in Section 2.5 on page 12, metals generally reflect far more at sharp incident angles than non-metals do. This is taken into account by the refractive index parameter. Some rendering systems choose to use a single refractive index value for metals, but increase it significantly. In doing so, the color-shift effect is lost.

4.1.6 *Microfacet reflection*

The specular part f_{spec} is computed for instance according to the GGX evaluation of Section 2.4.6 on page 12. It depends on a Fresnel term, which is either the dielectric or the conductor variant, and a microfacet normal distribution function (NDF). Facets are ideally specular. The NDF contains most of the computation.

4.1.7 *Microfacet transmission*

The microfacet model accounts for both reflection and transmission. It operates similar to the microfacet reflection BRDF but uses the sampled microfacet to refract on instead. It can be used to render rough transparent materials such as etched glass.

4.1.8 *Microfacet distributions*

NDFs have a significant impact on the end result. Two different NDFs are used: Beckmann and GGX (described in Section 2.4.4 on page 11). Images rendered in this document use the Beckmann NDF unless otherwise stated. While the GGX distribution has become increasingly popular, in some situations it produces considerable amounts of noise that take a long time to converge. The visual result of using GGX over Beckmann is longer tails on highlights and a closer fit to real-life measurements of glass in particular.

4.1.9 *Refractive layer*

Materials such as varnished wood, floors treated with epoxy, plastics and rubber, and various lacquered surfaces, all have distinct layers. In order to represent these types of surfaces a coating BxDF is introduced. This is defined in terms of the attributes of a dielectric layer to be on top of another BxDF. For instance, a plastic may be modeled with a diffuse base and a clear coating on top. As the coating obeys the Fresnel equations, the effects of effective roughness can be observed: grazing angles make these surfaces much more reflective. The coating may be smooth or rough. The implementation assumes that this BxDF is used alongside a reflective component - this BxDF models only refraction in the coating and the interaction with the base layer.

4.2 Materials

Materials can be defined in practically limitless combinations of BxDFs, but relatively few make sense in terms of energy conservation and practical use. The simplest way of defining a material is to use a single BxDF, such as the microfacet model described in section 2.4 on page 9, but many real-life surfaces exhibit more than one type of light interaction. Glass both reflects and transmits light, and a lacquered wooden floor or the clear coat layer on a car will reflect some light, and the remaining light will interact further with the layers underneath. Reflected light may also diminish once more on its way out due to the reflective and possibly absorptive effects of the lacquer. This calls for layered reflectance models to represent a variety of materials.

This section presents the material classes available to the user. For rendered images using these materials, see Section 5 on page 24.

4.2.1 *Layering and composition*

A natural way of representing layered materials is to use a tree structure of BxDFs: for glass in particular, one child would be specular and the other transmissive. During evaluation, the Fresnel term for glass (supplying its relative refractive index) would be evaluated, yielding the ratio between reflectance and transmittance. This value could also be used directly as the probability of choosing the direction to sample: either the reflected ray, or the transmitted ray. Both energy conservation and importance sampling are thereby straightforward. In general, a tree is an intuitive structure to use for representing materials consisting of two reflectance models, of which one is specular and the other either diffuse or transmissive, or possibly another similar composite node. Depth is likewise simple, as long as the leaves are not specular. This way of representing materials is also suitable when it suffices to describe a material as a linear mixture of two reflectance models, where each child is evaluated and linearly interpolated, but this representation quickly becomes unintuitive and impractical when constructing more complicated materials - the artist would have to manually set a blending function, which may not only likely break energy conservation but would also be difficult to produce well-behaved materials with.

A more practical way of representing materials is as arrays. It makes the system more flexible: a BxDF can be easily inserted into the stack, and as long as it is correct in isolation, it will be correct in composition; there are no blending functions at all (though a separate BxDF may well model this). Instead of blending functions, the components are weighted by their relative contribution during evaluation. The renderer can also more flexibly ignore certain types of BxDFs, whether for performance or as an artistic choice, while maintaining correctness, since all components are evaluated in a linear fashion. This is important for direct lighting, where some types never contribute to the result - that is, the Dirac delta functions that have singularities in their distribution.

When a ray intersects with a surface, the corresponding material of that object is determined. We choose to then evaluate all of the components of this BSDF. A cumulative distribution is meanwhile built from the individual probabilities of the components of that material. One of the components is then picked from the normalized cumulative distribution, such that more significantly contributing components are sampled more frequently. The sample is weighted accordingly, to ensure an unbiased result. This procedure makes it simple to extend the system with further BxDFs, as the end result will be correct as long as each of the components produces correct values in isolation.

Example: Glass material. Suppose that the cumulative distribution has $R = [0, \alpha)$ and $T = [\alpha, 1)$. The PDFs for each is exactly one, since both specular and transmissive distributions are Dirac delta functions. A pseudo-random sample $\xi \in [0, 1)$ selects within the distribution: if $\xi < \alpha$ we choose the reflected direction; otherwise we follow the refracted ray instead, to continue the light path. The resulting value of the BSDF is scaled by the probability distribution, such that less likely paths have larger contribution, but that more likely paths are sampled more frequently but have smaller weights.

A drawback of using a BxDF stack is that all components are evaluated despite the possibility of some components having very low contribution. In images with few samples this method performs better but is slower, and for images rendered over time it may be faster to avoid some components based on their probability. In practice, however, much of the computation is already done when the probability is known. For materials such as glass, one could evaluate the Fresnel function once and only then evaluate either reflection or refraction (with a probability equal to the ratio).

4.2.2 *Matte*

A simplistic matte, or *diffuse*, material can be achieved by using the Lambertian BRDF. If provided a roughness value that is greater than zero, the Oren-Nayar extension can be used instead to model rough, diffuse surfaces.

4.2.3 *Mirror*

This simple material utilizes one of the strengths in ray tracing: accurate reflection. There is only one BxDF associated with it - the ideal specular reflection. Note that the Fresnel equations are not used; the material simply reflects all incident light in the mirror direction. The reflectance parameter R attenuates the intensity and color of the incident light.

4.2.4 *Glass*

Glass is a typical dielectric material, and it reflects and transmits light. Light that is not reflected is assumed to either be completely or partly transmitted, with attenuation parameters R and T for reflectance and transmittance, respectively.

Glass is be modeled with the ideal dielectric reflection and transmission BxDFs, using the Fresnel equations to determine their ratios. To this end, the relative refractive index η must be provided. Air, for instance, has a refractive index value of roughly $\eta = 1$, while for glass it is roughly $\eta = 1.5$. Absorbance over the distance traveled within the medium can be specified with σ_i. The absorbance in the medium outside, σ_t, is normally zero.

The direction of the surface normal determines whether incident light is entering the medium or exiting it - this is accomplished by checking whether the normal and the incoming ray are on the same hemisphere. Once the Fresnel ratio has been determined, the resulting reflected and transmitted light follow the ideal cases, though the refracted ray changes direction depending on the relative refractive index.

The glass material actually models a generic dielectric that is transmissive - it can be used to model some gemstones as well, though double refraction and diffraction are not taken into account. The transmittance parameter allows for modeling the color tint found in for instance rubies, emeralds, and sapphires, though a better method may be to use the absorbance parameter. Examples can be found in Figure 6 on page 24.

Rough glass uses the microfacet reflection and transmission BxDFs to model various rough surfaces, e.g. anti-glare, etched, or ground glass, with a roughness parameter r.

4.2.5 *Metal*

Typical real-time shading models tend to do well in representing non-metals, but may struggle to accurately emulate the very different reflectance properties of metals. To achieve a metallic look, a single, one-dimensional refractive index value may be used - one that is much higher than non-metals to make it reflective. A complex index of refraction more accurately emulates the color shifts in grazing angles.

Since metals do not exhibit subsurface scattering, they are almost entirely specular, with the reflections dependent on both angle of incidence and wavelength. Some models make use of measured reflectance data for the relation between angle of incidence and the wavelength of the light. MERL [14] is an example of a database that provides measurements for common materials, but it is limited to research and academic use only. There is a relatively small additional computation cost associated with a larger number of samples.

Providing enough input for realistic metals remains a balance between practicality and visual quality. For this work, a judgment was made that one complex value (η, k) for each of the red, green, and blue wavelengths is sufficient. Values for specific wavelengths can be obtained from literature or online compendiums, such as [15].

A brushed metal appearance can be simulated by using anisotropic versions of the NDFs, where there is asymmetry in the tangential plane for the surface roughness. Texture maps, or tangents and bitangents for vertex data, may be used to set the orientation of the anisotropy. The parameter is the relation between the two dimensions $r_y/r/x$

Conductors do not transmit light, and reflectance is dependent on wavelength. The conductor specular reflection BRDF is used when roughness r is zero. Otherwise, the microfacet BRDF is used. These parameters are sufficient to produce accurate metals. Examples of metals are shown in Figure 12 on page 29.

4.2.6 *Plastic*

Plastic and rubber are dielectrics, which transmit light to some degree. This type of material is therefore more complicated than metal to render with accuracy and speed. In real-time contexts, a simplified model is often used that divides it into a diffuse part and a specular part, with factors for each, e.g. Blinn-Phong. For this implementation we choose to also divide plastics into a base part and a glossy-specular part. A microfacet model is used as the specular component, which for $r = 0$ reverts to the ideal version. For the base, we choose to use the refractive layer that contains a diffuse base (Lambertian or Oren-Nayar) to approximate the subsurface scattering effects that we are otherwise not computing. The coating layer is defined by its refractive index η, and the $\mathit{thickness}$ and $\mathit{absorbance}$ parameters which together affect how much absorption is simulated, and the reflectance, or base color, of the underlying layer R.

The result is that the base contains a layer that reflects and transmits incident light, bending the light path into the underlying surface, where it interacts with the base layer and reflects light outward. The reflected light interacts further with the dielectric layer, reflecting or transmitting light. The interactions between the base and the incident medium may occur an arbitrary number of times: light may reflect many times inside the coating. The relationship is a simple one - a geometric sum of reflections and transmissions - and the resulting value is adjusted to take into account this phenomenon. Otherwise there would be considerable energy loss for the material. A different attempt at dealing with energy loss for rough materials has also been suggested [16], but for smooth layers (non-microfacet models) this term cancels out, meaning there is still energy loss.

4.2.7 *Car paint*

Defined by a refractive layer with $\mathit{thickness}$ and $\mathit{absorbance}$; and a metal base of color c, roughness r, and per-wavelength refractive index (η, k).

Car paint behaves similarly to plastics. The base is metallic, with several layers added on top, including layers to prevent corrosion and a primer to level the surface. This material models metallic paint as a metallic core with a coating on top, in other words the base coat and the clear coat. The model is similar to the plastic material, except that the base layer is metallic instead of being matte.

Pearlescence and sparkle and other glitter effects in car paints are not modeled. Macro-level spatial variation is more appropriately taken into account with textures.

4.3 Fabrics

Various types of cloth is useful for representing for instance interiors of cars. Some fabrics are more difficult to model than others, such as velvet and satin which have strong back-reflecting (or retro-reflective) properties due to their grains lying down with a similar and particular orientation. A simplification is done to the matte material to weigh back-scattered light more heavily. The color is adjusted similarly by changing the reflectance parameter R, and the back-reflectance strength can also be adjusted.

5 RESULTS

Although materials are represented internally by BxDF stacks, they cannot be defined by the user directly. Instead, only prepackaged stacks are exposed in the form of materials, as presented in Section 4 on page 16. This simplifies the scene-setting process. Another benefit is that it encourages the use of physically-correct materials which are less error-prone when the camera or light sources are moved. These materials allow for simple tuning instead of having to understand the composition of particular surfaces and how they interact with light. This section demonstrates the shading model by way of renderings of the materials described.

All scenes are lit by a simple, uniformly white environment light, unless noted otherwise.

5.1 Cornell box

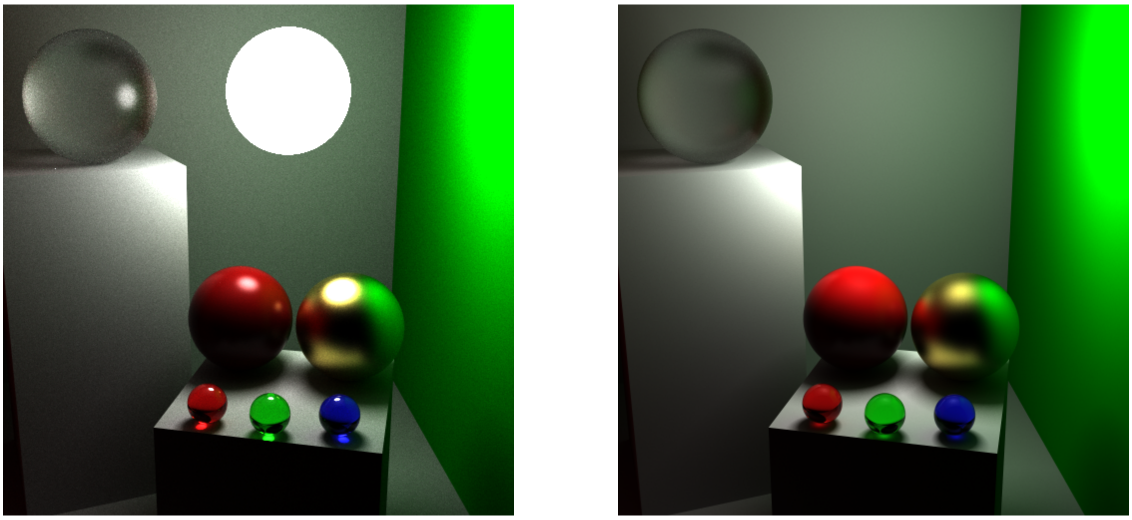

(a) Area light source. (b) Light source without geometry.

Figure 6: The Cornell Box scene (4096 spp).

A variation of the Cornell box, this scene (Figure 6) shows smooth matte surfaces on the boxes, ceiling, walls, and floor. A rough glass ($r = 0.1$) material is applied on the left-side sphere, and red plastic and rough gold ($r = 0.2$) are applied to the two larger spheres on the right. The frontmost small spheres are glass, with the refractive index and absorbance of ruby, emerald, and sapphire, demonstrating the range of the glass material to (singly refractive) gemstones.

The scene is lit by (6a) an emissive sphere and (6b) a spherical light without associated geometry. The relatively small light in (6a) causes slow convergence due to it contributing not only directly, but for light that happens to reflect off it. The second result shows the non-geometric light source which converges significantly faster. An equivalent result is achieved after roughly 100 samples in (6b), compared to the unconverged result in (6a) that has 4096 samples per pixel (spp). The specular lights in the second image are of the brightly lit ceiling, rather than the light source itself. One would expect a light source where the light originates, which is a problem with the second approach.

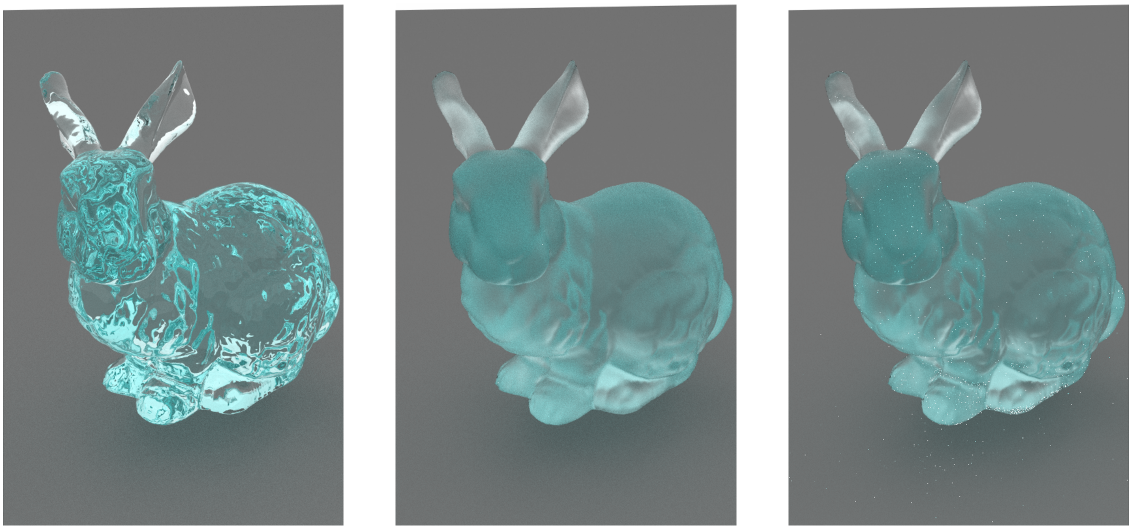

(a) $r = 0$. (b) $r = 0.2$ (Beckmann). (c) $r = 0.1$ (GGX).

Figure 7: The Stanford Bunny model rendered with absorptive glass of various roughness.

5.2 Stanford Bunny

The glass material can be used to model transparent, absorbing materials. The scene in Figure 7 shows an example of a solid glass object that absorbs only red light, with the amount based on the travel distance through the object. Less light can be seen to be absorbed through the thinner ears of the bunny.

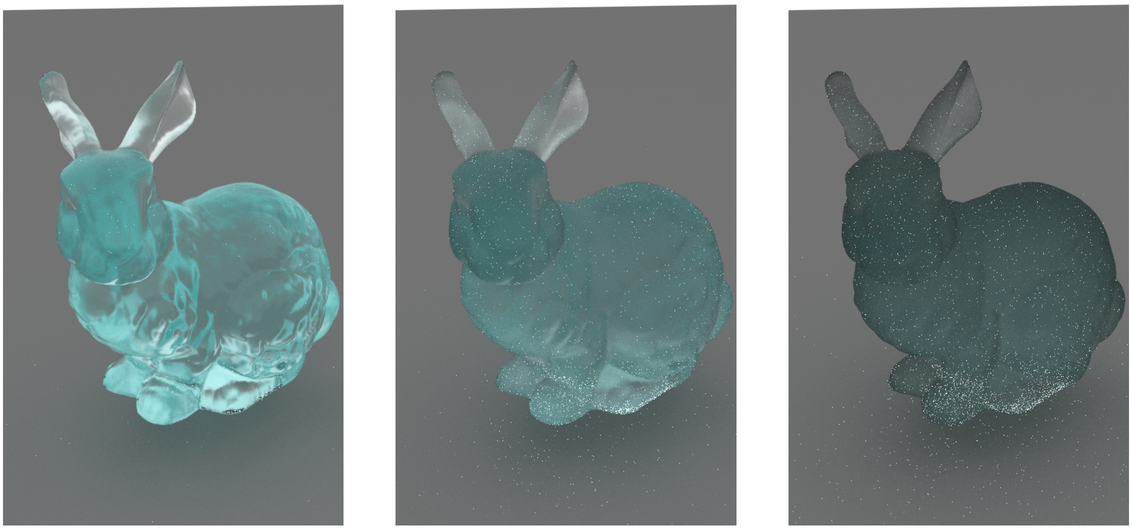

(a) $r = 0.05$. (b) $r = 0.2$. (c) $r = 0.4$.

Figure 8: Additional renders using GGX. A darkening effect occurs due to energy loss.

As seen in Figure 7c, the microfacet model causes significant noise in the form of fireflies (extremely bright pixels). This is particularly noticeable for high roughness values, and when using the GGX distribution.

(a) Matte smooth. **(b)** Matte rough ($r = 0.4$). **(c)** Metal rough (Alum.). **(d)** Metal rough (Gold).

Figure 9: The Happy Buddha model rendered with matte and metal materials. The metals have roughness $r = 0.2$. (c) reflects only red light, for the sake of comparison.

5.3 Happy Buddha

(a) Plastic ($r = 0.1$). **(b)** Ceramic ($r = 0$). **(c)** Car paint ($r = 0.005$). **(d)** Car paint ($r = 0.1$).

Figure 10: The Happy Buddha model rendered with coating of various roughness r. The base surface is aluminium, with a thick coating that absorbs mostly green and blue light ($\mathtt{thickness} = 1.8$, $\mathtt{absorbance} = (0.25, 1, 1)$, $\eta = 1.5$). The surface underneath the coating is (a) smooth, (b) 0.5, (c) 0.1, and (c) 0.2.

Figure 10 illustrates some coated materials. Plastics and ceramics build on the same material but have different parameter values: here, the ceramic has a smooth coating and a rough base, and vice versa for the plastic. In each case, the base surface is white or, in the case of metals paint, aluminium - the color is achieved only through the glazed coating, defined by its thickness, absorbance, and refractive index. In general, low thickness values lets more of the base color through, and a higher refractive index makes the coating more reflective. A higher base roughness yields duller color, which is particularly apparent on the head of the model.

(a) thickness $= 0$. **(b)** thickness $= 0.3$. **(c)** thickness $= 0.6$. **(d)** thickness $= 1$.

Figure 11: Shows the plastic material with varying thickness on the coating and a constant absorbance of $0.25, 1, 1$ in red, green, and blue, respectively. A thicker coating allows for more light to be absorbed, yielding a darkening effect. Since the coating absorbs each color to some degree, with high enough thickness, all color is absorbed and only the specular component remains. The underlying surface is gray ($R = 0.5, 0.5, 0.5$).

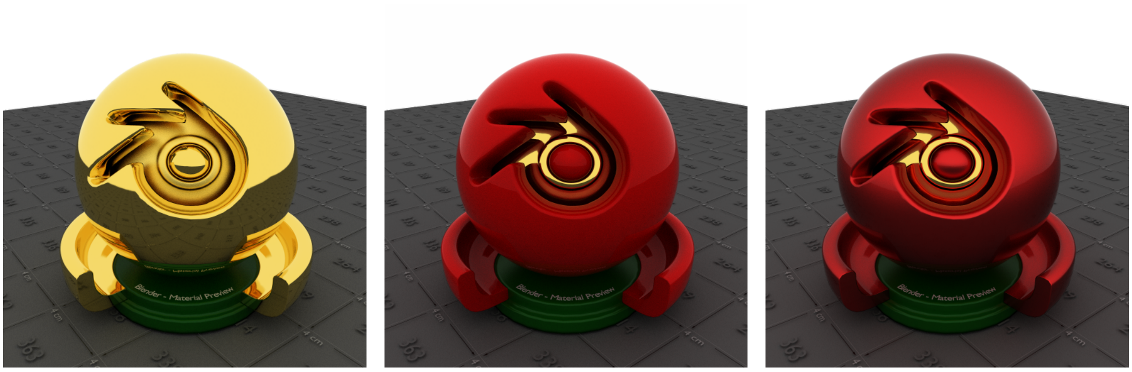

(a) Gold. (b) Plastic. (c) Metallic paint.

Figure 12: The Blender material test scene (author: Robin Marin).

5.4 Blender ball

Figures 12 and 13 on the following page demonstrate the materials applied to a different model.

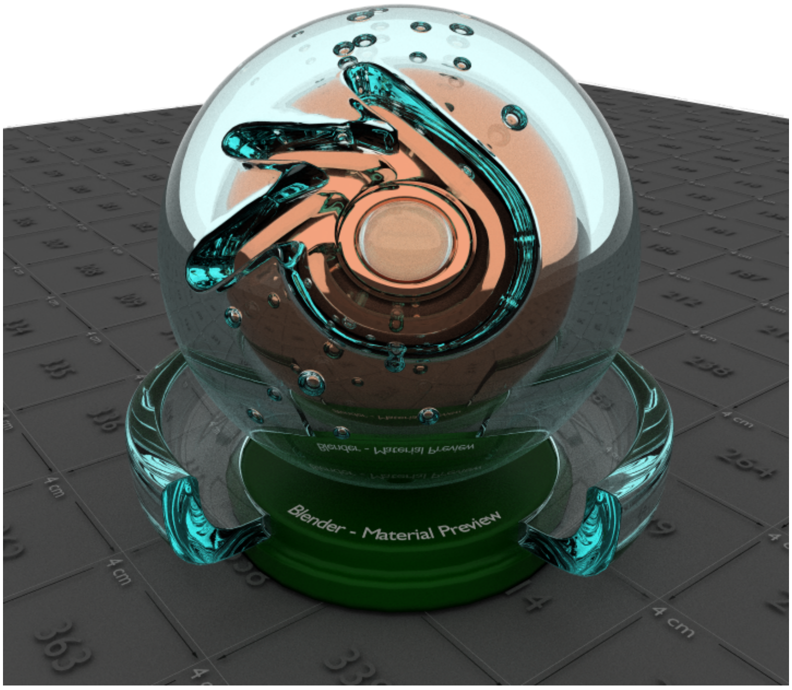

(a) Ambient light.

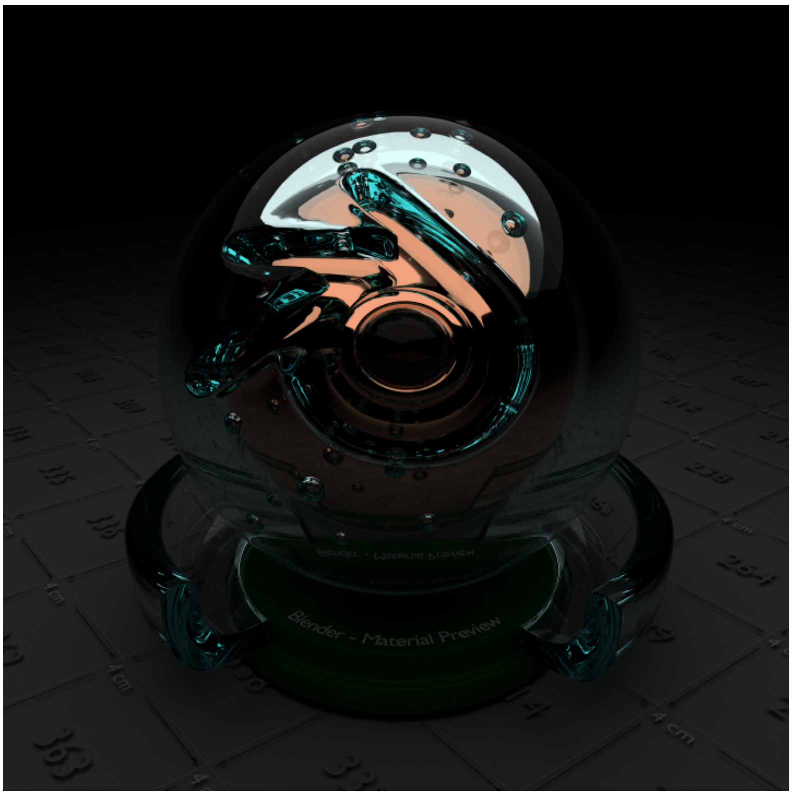

(b) Quad light.

Figure 13: Red-absorbing glass, polished copper, and at the base rough green plastic.

5.5 BMW

Figure 14 on the next page demonstrates the car paint material on a simple car model, along with various other materials. The glass is a single-face sheet of geometry, so a simplified glass material is used (not presented) along with the transmittance factor to achieve a tinted appearance. The rims are chrome and the tires use the plastic material.

(a) Ambient light.

(b) Quad light.

Figure 14: The BMW scene rendered with red-absorbing car paint with an aluminium core. The lighting is either from a uniformly white environment light (ambient) or a large quadrilateral light source above the car. (CC0)

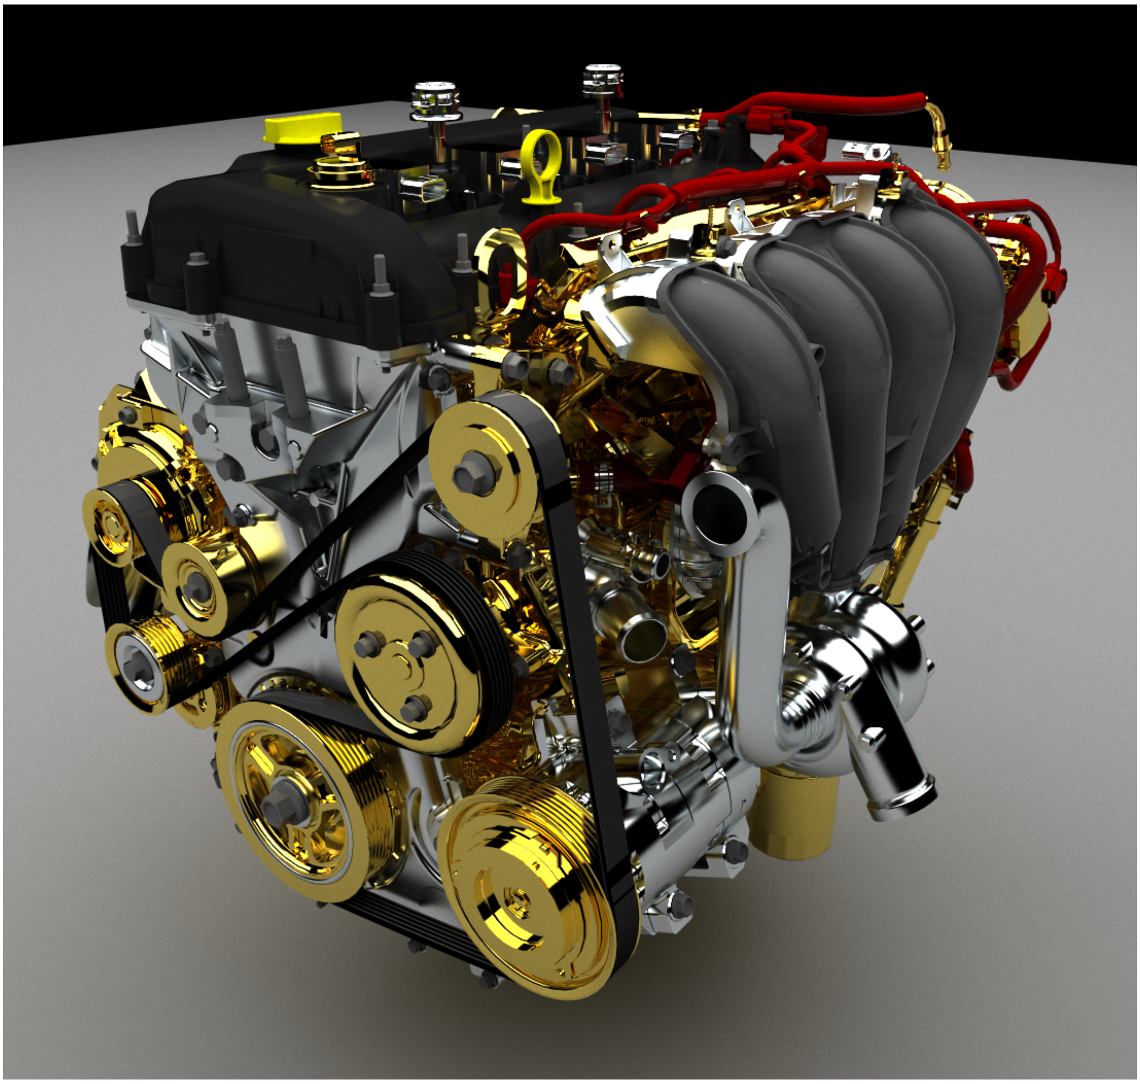

Figure 15: A complete engine consisting of 7.8M triangles.

5.6 Engine

Various materials are shown once more in Figure 15.

6 CONCLUSION

A simple shading model was presented that offers basic, representative materials within the domain of the automotive industry. The materials are the most commonly occurring ones, and are built through physically-based light interaction models.

Only simple names (e.g. metal) and parameters (e.g. reflectance, roughness) are exposed to users. Material attributes that are more difficult to reason about, whether in terms of obscurity (e.g. refractive index) or scale (e.g. absorption, roughness), are either optional or easily found in compendiums. Unintuitive parameters, such as the complex refractive indices of metals, can also be found in databases online and in literature.

A prototype renderer was implemented to demonstrate the shading model. It lacks many features of a complete rendering system beyond surface shading.

7 DISCUSSION AND FURTHER WORK

A darkening effect for rough glass was observed. In the original paper a non-physical scaling factor was used to account for energy loss, which is more difficult to put into context of the whole system.

The coating material also either introduces energy loss or has a discrepancy between the smooth and rough variants: it uses a different method to deal with energy loss in the two cases.

Although textures were not available for this project, they may be procedurally generated and projected on surfaces to allow for some spatial variation, including generating car paint glitter, or wood fibers with Perlin noise.

8 ETHICS

Photorealistic rendering may be used to deceive users about the existence of products before they are made.

REFERENCES

[1] Turner Whitted. An improved illumination model for shaded display. In *ACM SIGGRAPH Computer Graphics*, volume 13, page 14. ACM, 1979.

[2] 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.

[3] Matt Pharr and Greg Humphreys. *Physically Based Rendering, Second Edition: From Theory To Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2nd edition, 2010.

[4] Stephen H. Westin, Hongsong Li, and Kenneth E. Torrance. A field guide to BRDF models. Cornell University Program of Computer Graphics, January 2004. Technical report PCG-04-01.

[5] T. Akenine-Möller, E. Haines, and N. Hoffman. *Real-Time Rendering, Third Edition*. CRC Press, 2008.

[6] Eugene d'Eon Eric Heitz, Johannes Hanika and Carsten Dachsbacher. Multiple-scattering microfacet bsdfs with the smith model. In *Proceedings of ACM SIGGRAPH Asia*. ACM, 2016.

[7] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1112, 1967.

[8] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG)*, 1(1):7–24, 1982.

[9] Christophe Schlick. An inexpensive BRDF model for physically-based rendering. In *Computer graphics forum*, volume 13, pages 233–246. Wiley Online Library, 1994.

[10] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pages 195–206. Eurographics Association, 2007.

[11] Bruce G Smith. Geometrical shadowing of a random rough surface. *Antennas and Propagation, IEEE Transactions on*, 15(5):668–671, 1967.

[12] Ingo Wald, Sven Woop, Carsten Benthin, Gregory S. Johnson, and Manfred Ernst. Embree: A kernel framework for efficient cpu ray tracing. *ACM Trans. Graph.*, 33(4):143:1–143:8, July 2014.

[13] Michael Oren and Shree K Nayar. Generalization of lambert's reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 239–246. ACM, 1994.

[14] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769, July 2003.

[15] M. N. Polyanskiy. Refractive index database. http://refractiveindex.info, 2016. [Online; accessed 27-april-2016].

[16] Andrea Weidlich and Alexander Wilkie. Arbitrarily layered micro-facet surfaces. In *Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and Southeast Asia*, pages 171–178. ACM, 2007.