Arbitrarily Layered Micro-Facet Surfaces

Andrea Weidlich and Alexander Wilkie

Abstract

In this paper we present a method to combine several micro-facet based surface layers into a single unified, expressive BRDF model that is easy to use. The restriction to micro-facet based layers constitutes no loss of generality, since both perfectly specular and perfectly diffuse surfaces can be seen as limit cases of the micro-facet approach.

Such multi-layered surfaces can be used to re-create the appearance of a wide range of different materials, and yield good results without having to perform explicit sub–surface scattering computations.

This is achieved through suitable approximations and simplifications of the scattering within the simulated layered surface, while still taking absorption and total internal reflection into account. We also discuss the corresponding probability distribution function that is needed for sampling purposes, and investigate how the flexibility of this new approach is best put to use.

CR Categories: I.3.7 [Computer Graphics]: Three–Dimensional Graphics and Realism

Keywords: surface model, global illumination, microfacets

1 Introduction

During the last decades a number of analytical reflectance models have been developed for computer graphics use. These models can roughly be divided into two groups:

1. *empirical* – and usually physically implausible – models, which deliver reasonably good-looking results at moderate computational cost, and
2. those where comparatively expensive, *physically based* computations of light interacting with matter are used for highly convincing depictions of surfaces.

So far, complex layered surface models were not in general use in photorealistic rendering, mainly because the derivation of a compound BRDF is rather difficult for a general arrangement of layers. However, under the – not unrealistic – assumption that the layers involved are thin compared to the size of the micro-facets, the problem can be simplified in a way that still yields plausible results, but maintains the flexibility of a full layered surface simulation.

The goal of this paper is to present a simple but still physically plausible BRDF model that can simulate both smooth and rough multi-layered surfaces, and which includes absorption within the layers, as well as total internal reflection. Our new model offers the user a large flexibility in terms of what surfaces can be described with it, while still being intuitive to use, and – perhaps most importantly – also still physically plausible, despite the simplifications we apply to the problem.

The paper is structured as follows: we first give a short overview of related work. In the main part of the paper we present our multi-layered surface model, and conclude by validating our results, and giving an outlook to future research areas.

Figure 1: The surfaces of this art exhibit are all modelled with our BRDF approach. Apart from the bellow, the figure is entirely made of aluminium, but some parts have been coated in various ways using the methods discussed in this paper. The colour of the original material can be seen on the buttons.

2 Background and Related Work

The prior research relevant to this paper falls into three main categories: micro-facet based reflectance models, multi-layer reflectance models (as the two base techniques from which our proposed method is built), and Monte Carlo image synthesis.

For the purposes of this paper, the latter is only relevant insofar as we have to briefly discuss the requirements of modern rendering algorithms with respect to reflectance models, in order to describe our proposed model in a way that is actually useful in a real global illumination renderer.

2.1 Micro-facet Reflectance Models

Models based on a micro-facet approach are normally used to simulate rough surfaces. They assume the surface to consist of a large number of very small statistically distributed micro-facets, which are oriented according to some given probability distribution function, and which can be either isotropic or anisotropic.

For the purposes of this model, one considers a surface to be a collection of a large number of tiny, symmetric V cavities with two opposing facets. These facets are assumed to be perfect mirrors, the reflectance of which is governed by the Fresnel terms. The model takes mutual masking and shadowing between the facets into account. Within the constraints of the model, the specular reflectance which results from this approach is physically accurate, and given by

$$f_r = \frac{FDG}{4 \cdot (N \cdot L)(N \cdot V)} \tag{1}$$

- D is the distribution function of the micro-facets.
- G is the geometric attenuation term that influences self-shadowing when the incident light is blocked, and self-masking when the reflected ray is blocked.
- F is the Fresnel term for each micro-facet which describes the amount of light that is refracted and reflected.

Micro-facet theory was brought to computer graphics proper in 1982 by Cook and Torrance[Cook and Torrance 1982], who introduced a somewhat simplified and refined version of the original Torrance-Sparrow model that had been adapted to graphics use. Other micro-facet surface models are those of Oren and Nayar[Oren and Nayar 1994] (which is insofar unique, as it is up to now the only technique to use Lambertian micro-facets), the Ward model[Ward 1992] and the model of Ashikhmin [Ashikhmin and Shirley 2000].

2.2 Multi-Layer Reflectance Models

Layered surface models offer a great potential for creating very convincing renderings, and have already received a considerable amount of attention in computer graphics. The following list is by no means complete; due to space limitations we only give a brief overview of those papers which are most relevant to our work.

Classical layered surface models are those of Kubelka and Munk [Kubelka and Munk 1931] and Hanrahan and Krueger [Hanrahan and Krueger 1993]. Both these models have no closed mathematical form, and are therefore rarely used in practice – at least for image synthesis purposes. Based on the Kubelka–Munk model, Dorsey and Hanrahan [Dorsey and Hanrahan 1996] used layered surfaces to achieve surface aging effects. Icart and Arques [Icart and Arquès 2000] [Icart and Arquès 1999], as well as Hirayama et al. [Hirayama et al. 2001a] [Hirayama et al. 2001b] [Hirayama et al. 2000], presented approaches to accurately calculate the reflectance properties - including interference effects - of multilayer films. Another model that describes interference and dispersion is that of Granier and Heidrich [Granier and Heidrich 2003] .

Neumann and Neumann [Neumann and Neumann 1989] were amongst the first to propose layered surface models. They discuss two models, one that consists of a single perfectly smooth, transparent layer over an arbitrary surface, and one with an arbitrary number of layers. Both these models include absorption, but not internal reflection. Since they do not give a sampling PDF for their compound BRDFs, and neither a closed expression for the entire BRDF nor an algorithm to compute it, their work has to be considered somewhat incomplete.

Kelemen and Szirmay–Kalos [Kelemen and Szirmay-Kalos 2001] used the Cook-Torrance model in conjunction with layered surfaces, albeit in a simplified form. Their model lacks the ability to simulate absorption and internal reflections, and relies on a simplified variant of the Cook-Torrance model for its specular component. The diffuse component is described by a Lambertian term, and the combination of the two is dependent on the incident angle. They also provide an efficient scheme for sampling the BRDF in a stochastic renderer.

Wilkie et al. [Wilkie et al. 2006] use a multi-layer reflectance model with a transparent, rough dielectric layer over a normal diffuse surface to describe the reflectance properties of diffuse fluorescent surfaces, such as cardboard.

Another relevant surface model is that of Schlick [Schlick 1993]. Although it is not an actual layered model, it should be mentioned here since it can be used to mimic the appearance of layered surfaces fairly well.

The appearance of multi-layered models was also indirectly simulated by Lafortune et al. [Lafortune et al. 1997]. A wide number of different surfaces can be reproduced by their technique, but it has two disadvantages: there is a high discrepancy to real surface behaviour near grazing angles, and it is a purely empirical approach.

2.3 Monte Carlo Image Synthesis

While images of perfectly diffuse surfaces and perfect mirrors can be rendered by a deterministic ray-tracer, surfaces with arbitrary reflection properties are basically only tractable through Monte Carlo rendering, e.g. bi-directional path tracing or photon tracing.

2.3.1 Sampling PDFs

Common to such expansion solvers of the rendering equation is the fact that at each recursion level, they attempt to evaluate the illumination integral through stochastic numerical integration. One of the standard techniques to accelerate the convergence of such a stochastic integration is to perform *importance sampling*, which requires that the integrand be randomly sampled using a probability density function that mimics the integrand as closely as possible.

In practice, this means that for any reflectance model one not only needs formulas for its BRDF values, but also an efficient sampling PDF for the BRDF. While formulas for the BRDF are usually given in literature, a sampling PDF is often omitted, which limits the immediate applicability of some published models.

2.3.2 Path Propagation vs. BRDF evaluation

It is rarely explicitly mentioned in rendering literature that one has to be able to perform two distinct functions for each reflectance model one wishes to include in a stochastic renderer.

1. The first concerns the ability to correctly continue an incoming path according to a chosen sampling PDF. This is rather easy to perform even for complicated multi-layer surfaces: for each layer a suitable propagation direction can be recursively calculated, and is weighted according to its sampling probability. One of these rays is then followed by random selection.
2. However, one also has to be able to evaluate the *entire*, *combined* BRDF for arbitrary input and output directions. For arbitrary multi-layer surfaces the computation of this second piece of information is far from being trivial.

3 Arbitrarily Layered Micro-Facet Surfaces

As noted earlier, the idea of using individual surfaces of somewhat limited applicability – such as perfect mirrors or Lambertian surfaces – as layered components of a more sophisticated BRDF, is immediately appealing due to its simplicity and usefulness.

While the concept of using layered surfaces is simple, actually using it in a renderer is not – at least if the unrestricted case is considered, because in this case the computation of the entire BRDF (item 2 from section 2.3.2), would – if done correctly – involve sub-surface scattering computations within the layers.

3.1 Simplification of the Problem

The key to using such surfaces without expensive sub-surface scattering computations is to perform four simplifications:

1. Any micro-facets are considered to be much larger in horizontal extent, than the layers are thick.

2. All rays that are generated by sampling of lower BRDF layers, are assumed to exit at the original point of incidence.

3. Refraction rays that are generated for the computation of the entire BRDF, are assumed to meet at a single point on the next layer interface.

4. All light scattering is due to reflection at the boundaries between layers; no scattering occurs within individual layers.

See figure 2 for a sketch of the simplified reflection and ray propagation geometry, and figure 4 for a BRDF evaluation sketch. A detailed description of how exactly simplifications 2 and 3 turn out to be useful is given in section 3.2.

It is worth noting that none of these assumptions is entirely implausible: the first is consistent with the notion of only applying very thin layers atop a base substrate, while the second and third are reasonable simplifications under the circumstances – especially when one considers that the micro-facets involved are assumed to be statistically distributed entities in the first place.

In addition to these three simplifications, we impose one additional restriction: namely that the material used for any (partially) transparent layers only attenuates light passing through it, and does not contribute any secondary scattering effects of its own. Since most clear and tinted varnishes (which are one of the main targets for this model) do not exhibit noticeable scattering, this is not a particularly hard restriction, though. We also assume any varnish layers to be homogeneous. The three simplifications, taken together with the restriction to non-scattering varnishes, are what allows us to omit a full sub-surface scattering computation.

It has to be noted, though, that due to the simplifications we perform, our new approach does not constitute a general solution for surface layers of arbitrary thickness and is limited to surfaces that have layers thick enough to have the influence of absorption, and sufficiently thin enough that our simplifications hold true. However, given the high quality of the results obtained by our method, we do not see this as an immediate concern, at least not for those layered surface types which are readily modelled by our approximative approach.

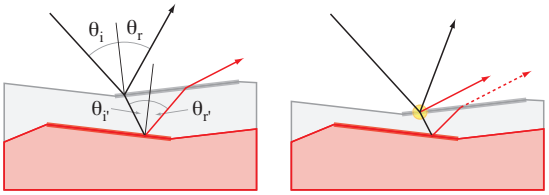

Figure 2: Computation geometry and simplification of the sub-surface scattering in a layered surface during BRDF sampling. Any micro-facet is large in relation to the layer thickness, which allows us to do the following: **(1)** a ray will always leave through the same micro-facet that it entered through, thereby eliminating the need to perform an intersection test with the nearby micro-facets. And **(2)**, any exitant ray coming from a lower level will emanate from the original point of entry (yellow dot), regardless of how oblique the exitant angle is. The direction of these rays is computed according to the correct geometry (dashed line).

3.2 Overview of the Model

Using our layered BRDF model poses two separate problems, that correspond to the two tasks outlined in section 2.3.2:

1. We have to be able to cast samples according to the BRDF.

2. We have to be able to compute the entire BRDF for arbitrary input and output directions.

The first problem can be solved by a recursive process that works as follows:

1. Any light that hits an interface in the layer stack is partly reflected, and partly refracted. The actual amount of energy that will be reflected is determined through Fresnel reflectance calculations – any surface in the stack (except for the very last one, where this is optional) is assumed to be governed by the Fresnel terms at least on a micro–facet level. An appropriate sampling direction is generated for the reflective component.

2. The refracted part of the energy is assumed to enter the material. A part of this will possibly be absorbed by the varnish material, and the rest then interacts with the next, second surface in the stack - the process is recursively started at step 1 for this interface.

3. All light that is reflected from lower layers is again attenuated by the varnish on its upward path, and possibly subjected to total internal reflection. This means that any directional samples from lower layers have to be treated accordingly during the return from the recursion.

Figure 2 gives an overview of the geometrical simplification associated with this process of generating directional samples. The second problem – computation of the entire BRDF – also requires a recursive approach:

1. The BRDF of the topmost level f_{r_1} is evaluated for the two given, arbitrary incoming directions ω_i, and ω_o. This yields a reflection component, and, except at the lowest layer, two refraction directions.

2. Any energy that is refracted into the next level T_{12} follows the two refraction directions associated with the initial incident directions, and is partly absorbed a by the medium.

3. These two refraction directions are assumed to meet at a single point on the next layer f_{r_2}, and the process is repeated from step 1 until an opaque layer without a refraction component is encountered.

4. On returning from the recursion, the individual BRDF components are attenuated by the Fresnel transmission coefficients T_{21} for the level above them, and added to the total BRDF.

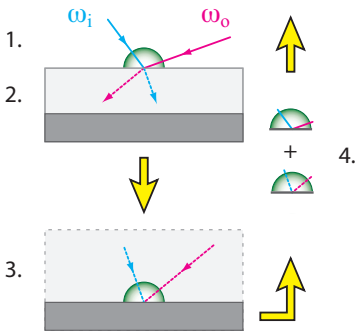

Figure 4: The recursive BRDF evaluation process used by our method; the numbers correspond to the steps described in the text.

Figure 4 illustrates the concept, and mathematically this also can be summarised as

$$f_r = f_{r_1}(\theta_i, \theta_r) + T_{12} \cdot f_{r_2}(\theta_{i'}, \theta_{r'}) \cdot a \cdot t \qquad (2)$$

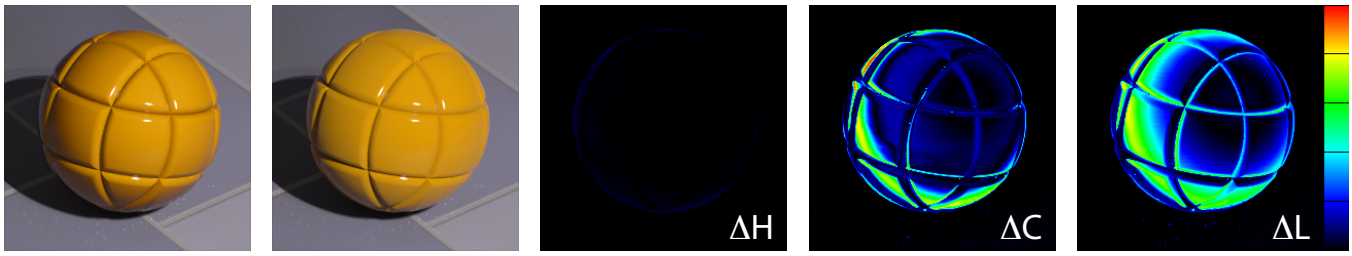

Figure 3: A surface of type **b)** (as defined in figure 6) rendered with and without absorption, with otherwise identical colour, illumination and surface parameters. The three difference images are – from left to right – for the hue, chroma and lightness channels of these two images. As expected, there are practically no hue differences, but chroma and lightness change considerably, mainly at grazing angles.

with

$$a = e^{-\alpha d \cdot (\frac{1}{\theta_{i'}} + \frac{1}{\theta_{r'}})}. \tag{3}$$

and

$$t = (1 - G) + T_{21} \cdot G \tag{4}$$

Note that the Fresnel reflection is already included in $f_{r_1}(\theta_i, \theta_r)$ in formula 2.

3.2.1 Absorption Term

Usually a part of light is absorbed inside a transparent material by travelling through it. The intensity of the absorption is defined by the Bouguer-Lambert-Beer law. According to that, the loss of intensity I of a light wave travelling through a material with a thickness of l is related to its initial intensity I_0 through

$$I = I_0 e^{-\alpha l} \tag{5}$$

where the constant α is the wavelength dependent absorption coefficient which defines the extent to which a material absorbs energy.

Note that the length of the path is determined by the thickness of the layer d as well as the cosines of the incident angle $\theta_{i'}$ and the outgoing angle $\theta_{r'}$ so that

$$l = d \cdot (\frac{1}{\cos\theta_{i'}} + \frac{1}{\cos\theta_{r'}}). \tag{6}$$

According to this law, internal absorbance, and thus the corresponding colour, increases with l. As a consequence the colour decreases and the saturation increases, but additionally it can change the hue. This phenomenon can be seen in action in figure 5, and figure 3 shows a comparison between the same surface being rendered with and without absorption taken into account.

3.2.2 Total Reflection Term

When light propagates from a denser medium into a less dense one, and the angle of incidence is greater than the so-called critical angle, no light enters the second medium; this phenomenon is called total internal reflection (TIR). If the index of refraction of a medium is very high, a light wave can easily become trapped inside a medium, and total internal reflection may occur several times before even a part of the light wave is able to leave the medium. The remaining part is reflected again and emerges somewhere else. An exact computation of the effect is problematic, since it would require an explicit simulation of the incident ray interacting with the micro-facet geometry multiple times. Any practical solutions are approximations of varying accuracy.

As with the entire BRDF, adding a TIR approximation to the layered BRDF model poses two separate problems:

1. It has to be taken into consideration when casting a ray upwards from a lower level during BRDF sampling, and
2. it has to be included when computing the entire BRDF for arbitrary input and output directions.

The first is comparatively easy: sampling directions created on lower levels are culled if they would be affected by TIR. This effectively causes an energy loss, which we compensate by adding a term, that accounts for the energy that is not able to immediately leave the medium. Conceivably, one could also re-start the downward propagation, although it remains to be seen whether the difference to the results obtained with the energy compensation term would be worth the additional computational effort.

The second task is also not hard, but requires a "backwards" approach that is perhaps not immediately obvious: for a given incident direction, it is easily possible to determine which reflection direction in the lower medium could have caused this particular exit direction. The energy is then attenuated by the Fresnel transmission coefficients, and again a compensation term is added.

4 Sampling PDFs for Multi-Layered Surfaces

4.1 Micro-Facet Distribution Functions

Since our BRDF consists of many different functions, it is impractical to find a probability distribution that matches the whole BRDF. Instead, it is a common method to sample the component which most influences the appearance of rough surfaces, and use this distribution function to obtain the probability that a micro-facet is oriented in a specific orientation. In literature, several different probability distributions exist to predict the distribution of micro-facets; we decided to use the normalised Blinn distribution

$$D(\omega_h) = \frac{m+2}{2\pi}(\omega_h \cdot n)^m \tag{7}$$

where ω_h is the half angle and m is a measure of the roughness. The larger the average slope m of the micro-facets is, the more the reflection is spread out.

4.2 Sampling PDFs for Multi-Layer Surfaces

As discussed in section 2.3.1, it is crucial to find a good sampling PDF for any surface model used with MC image synthesis. Again we have to distinguish between the two tasks of section 2.3.2. The first task is rather easy. The components of the surface can be sampled individually according to their respective PDFs. As a consequence, importance sampling is still possible although a surface consists of BRDFs with different shapes. Which ray is followed is selected randomly.

Figure 5: Diffuse white spheres with a yellow varnish layer of varying thickness. Layer thickness values are 0.0, 0.4, 0.8, 1.6, 3.0, 5.0 and 15.0. Note the progressive changes in colour, saturation and hue.

The second task is a little bit more complicated. What we want is a PDF that is as proportional to our BRDF as possible, and of course fulfils the two required properties that negative probabilities do not exist, and that the sum of the probabilities of all possible outcomes over a domain D is 1.

$$p(x) \geq 0$$

$$\int_D p(x) d\mu(x) = 1$$

To achieve this we can weight the probabilities p of the individual PDFs of the compound BRDF with constant values w so that

$$p = \sum w_i p_i \tag{8}$$

If both BRDFs have a similar shape (e.g. two Torrance-Sparrow surfaces of equal roughness), the weighting factors are equal. Otherwise we have to choose a bigger weight for the BRDF that dominates the appearance of the whole BRDF. The respective value depends on the reflection properties of the different layers. However, the sum of these weights has to be 1.

5 Differences to Existing Techniques

The considerable body of related work in this area – which we outline in section 2 – mandates a detailed discussion of the novelty factor of our proposed method:

Combination The multi-lobe approach as suggested by Cook and Torrance works well for some surfaces (see [Ngan et al. 2005]), but cannot be used to describe retro-reflective surfaces, because a backscattering lobe cannot be expressed in a physically correct way by combining several forward scattering lobes. Also, their multi-lobe approach does not work for strongly dissimilar lobes, e.g. a combination of Oren-Nayar and a highly specular lobe. We instead opted to combine suitable BRDFs of arbitrary type – which can include the genuinely retro-reflective Oren-Nayar model if this property is needed – in a physically correct way, and as a consequence can produce a much wider range of BRDFs, than the combination of quite similar forward-scattering lobes could achieve on its own.

Patinas Our technique can be seen as being complementary to, rather than an improvement on, the work of Dorsey and Hanrahan [Dorsey and Hanrahan 1996]. A large part of their work deals with determining where layers of patina should be placed on a model, and which layers are present in a given patina. Conceivably, one could use our technique to evaluate the resulting compound BRDFs in a GI renderer as an alternative to the method proposed by the original authors.

Weighting We do not use a fixed parameter to determine the ratio between diffuse and specular components of the reflected light, like e.g. the classical Cook-Torrance model does. Since in reality the specularity of a surface increases with decreasing incident angle, such a fixed ratio between these components is physically implausible. Like Ashikhmin and Shirley [Ashikhmin and Shirley 2000], we use the Fresnel coefficients as a weighting factor instead.

Physical Correctness In contrast to the empirical Lafortune model (which just combines Phong lobes to fit a given BRDF at given angles of incidence), our model combines its components in a physically based fashion through the simulation of multiple surface layers. Also, we go further than Kelemen and Szirmay–Kalos [Kelemen and Szirmay-Kalos 2001], since we use the original, non-simplified versions of the BRDFs we combine.

Absorption So far, the only ones to investigate absorption within a BRDF model layers were Neumann and Neumann [Neumann and Neumann 1989], although they only discuss the case of a smooth surface over a Lambertian base substrate. Our model is not limited to that, and can handle absorption between layers of arbitrary roughness.

Inter-reflections We simulate the internal reflection that occurs when light travels from a denser to a less dense medium. As far as we know, nobody has done this before for a layered BRDF model.

Importance Sampling As mentioned earlier, it is crucial in the context of Monte Carlo image synthesis to be able to perform proper importance sampling on any reflectance model. Like [Kelemen and Szirmay-Kalos 2001], we provide a matching sampling PDF for our BRDF, even though our model is potentially considerably more complex.

Simplicity The main difference of our technique to more complex surface models such as e.g. [Icart and Arquès 2000] or [Hanrahan and Krueger 1993] is that it is not intended to be a comprehensive simulation of truly arbitrarily layered surfaces. Instead, we focused on providing a much simpler technique, that, with a few simplifying assumptions, still offers most of the benefits one can achieve with complex models, as well as being more intuitive and flexible to use.

6 Results

We implemented the BRDF proposed in sections 3 and 4 in a spectral rendering research system that features unbiased global illumination rendering algorithms, and tested the new reflectance model in various configurations.

The two big advantages of our approach are that it is physically based (if not entirely physically *correct* in the narrow sense of the word because of the approximations we introduce into the evaluation), and simple and intuitive to use. As can be seen in the next two sections, it can be used to simulate a wide variety of real surfaces.

We variously used the original Torrance-Sparrow model, as well as the Oren-Nayar and Lambert models as basic components f_{r_n} in

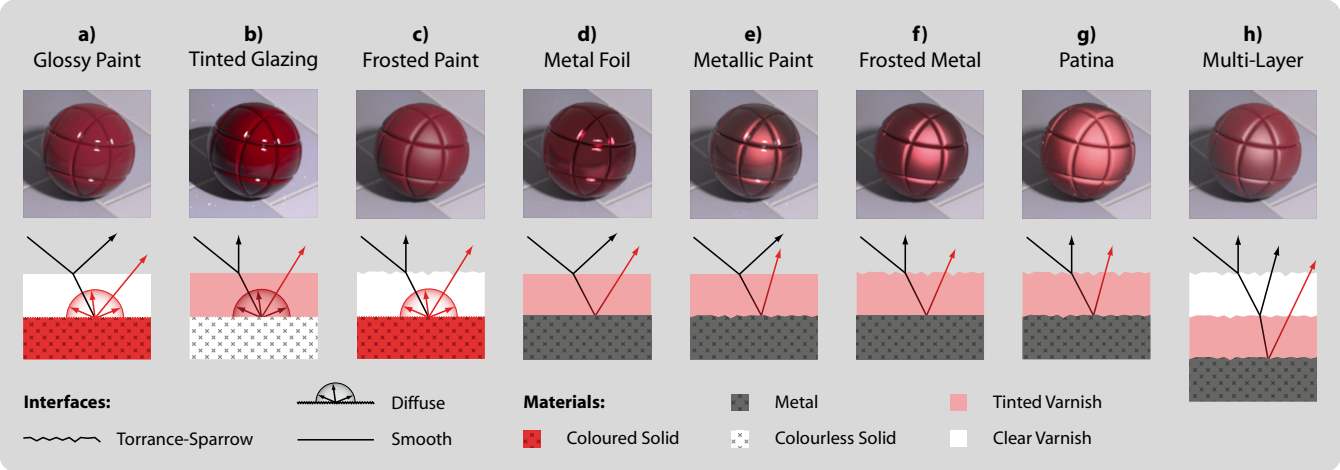

Figure 6: Examples of various surface types that can be generated by using our layered model in different configurations. The spheres are all rendered under the same illumination, and correspond to the cases discussed in section 6.1. It is important to note, that for reasons of clarity the surface layers shown here do not exhibit the simplifying assumption shown in figure 2 which we use for all our actual BRDF computations. In order to properly distinguish the various cases, the micro-facets are much smaller than the layer thickness in this drawing. Also note that the colour of the various varnish layers in this picture is only indicative of typical usage, and that more combinations than are shown here are possible.

our layered BRDF. So far we have not used any of the more recent analytical models, but not because they could not easily be included in this scheme. Our restriction to these three models was mainly due to the fact that we wanted to explore the behaviour of layered BRDFs with just a few representative and well understood base BRDFs used as building blocks, and these three proved entirely sufficient for this purpose.

6.1 Capabilities of the Proposed Model

The surface types discussed here correspond to the examples shown in figure 6, and it should be noted that this is probably not an exhaustive list of what can be achieved by this technique. Three of the examples were inspired by BRDF measurements from Cornell [COR]; for these scenes we attempted to match the appearance of the original images as closely as possible.

All these different surfaces were simulated using the same surface model code, with only the layering and the parameters of the individual interfaces and varnish layers being changed.

Surface type a) – Glossy Paint A typical application for a Lambertian surface covered with a smooth layer is the simulation of glossy paint or opaque ceramic glazing. Examples can be seen in figures 8a, where the clear varnish on the crayons has an index of refraction of 1.3, while the refractive index of the clear coating on the mug is 1.6. This is realistic, since ceramic glazing has a higher index of refraction than the binder typically used for enamels – as a consequence, ceramic glazing has brighter specular reflections than paintwork.

Surface type b) – Tinted Glazing This type simulates a (usually colourless) Lambertian surface covered with a smooth layer of tinted varnish. The result resembles certain types of transparent ceramic glazing, and examples can be seen in figure 5. For surfaces of this type, the colour usually increases in saturation towards grazing angles – compare this to the behaviour of case **a)**, where the colour largely stays the same everywhere on the object regardless of curvature.

Surface type c) – Frosted Paint Figure 7a shows a red sphere that has been coated with a gloss-reducing finish; this is one of the examples from the Cornell BRDF database [COR]. We combined a red Lambertian surface with a rather rough Torrance-Sparrow varnish. The average micro-facet slope is 12° and the refractive index of the varnish is 1.6. Here, the varnish layer is rather clear and thin – the thickness is 0.5.

Surface type d) – Metal Foil The combination of two glossy layers can be used to simulate surfaces that resemble colourful metallic foil. This effect can be reproduced by coating a smooth metallic surface with a smooth, tinted dielectric layer; the orb in the background of figure 8b is a good example of this. To create its surface, we took a silver sphere and coated it with a green varnish. Both layers have a surface roughness of 1°.

Surface type e) – Metallic Car Paint A nice example of the capability of our approach is the automotive paint from figure 7b, which we manage to simulate with only one layer over a metallic substrate. The base layer is a rather rough aluminium surface, with an average micro-facet slope of 12°. The upper layer simulates the varnish layer that coats the rough underlying surface, and is a very smooth transparent Torrance-Sparrow surface with roughness of 0.1° and an index of refraction of 1.45. The colour of the paint is only due to absorption in the tinted varnish. It is no coincidence that this yields good results, since this model is very close to the physical properties of real metallic car paint.

Surface type f) – Frosted Metal Frosted metal surfaces – such as those on the matte christmas orbs that can be seen in the foreground of figure 8b – can be created by combining a rough dielectric layer with a smooth metallic base surface. Here again a silver sphere served as the base, and was coated with a tinted dielectric layer with an average micro-facet slope of 12°. Refractive index and layer thickness are as in case **c)**.

Surface type g) – Patina Metals tend to oxidise when exposed to air. This effect can be simulated with a combination of the orig-

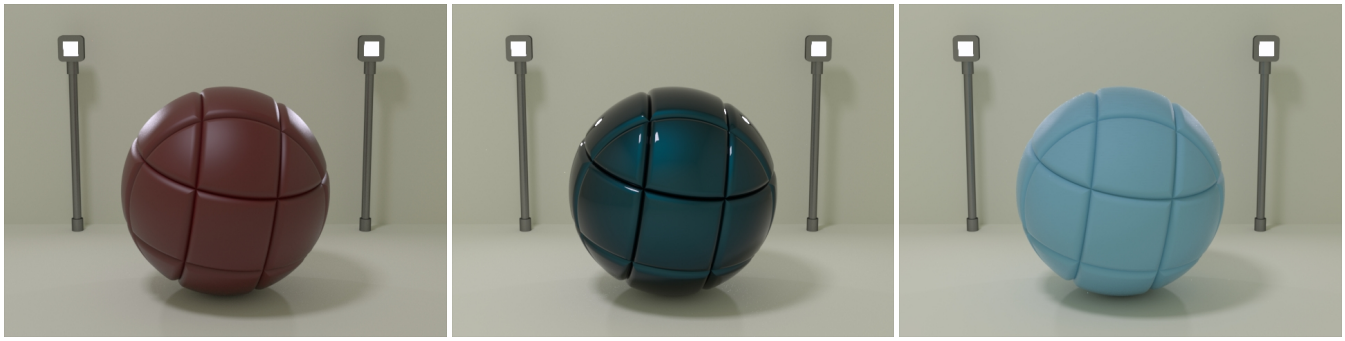

(a) (b) (c)

Figure 7: Left: Sample of red spray paint coated with glass frosting spray. This is an example of surface type **c)** from section 6.1 and figure 6. Middle: Dupont Cayman car paint sample. This is an example of surface type **e)** from section 6.1 and figure 6. Right: Blue wall paint sample. This is an example of surface type **h)** from section 6.2 and figure 6. The surface parameters for this image were obtained only through visual matching, since this particular BRDF data set at [COR]is incomplete – parts of the specular lobe are missing.

inal rough metallic surface, and an oxide-coloured rough dielectric varnish layer; figure 8c shows an example of this. The vase itself is made of pure, clean but slightly rough copper (which has a bright colour, as can be seen on the clean vase in the background) and has an average micro-facet slope of 8°. The vase in the foreground is coated with a strongly absorptive rough dielectric layer – roughness 12° – with weak reflectivity ($n = 1.3$).

Surface type h) – Multilayered Varnish For most applications it is apparently sufficient to use only one varnish layer. However, some surfaces have a microfacet distribution that cannot be simulated with only one layer. An example of such a surface is discussed in section 6.2.

6.2 Comparison with BRDF Measurements

For one of the BRDFs from the Cornell database [COR], we provide a comparison of both our model, and – for sake of completeness – the performance of the Lafortune compound BRDF [Lafortune et al. 1997] to measured reflectance data. We chose this particular surface because the measurement data were already interpolated and plots are available.

In figure 9, we show BRDF plots of three representative angles for the original data, our BRDF, and the Lafortune approximation we obtained with the published coefficients. Which, even though we used the same coefficients, is surprisingly not identical to the results in their paper; we assume that the authors plotted the BRDF for a different angle of incidence than stated in the plot labels.

In this case, we determined the composition and parameters for the layered BRDF manually through trial and error, because we did not have an automatic fitting algorithm. However, there is no reason to assume that an automated technique for the automatic fitting of BRDFs with our approach could not be developed.

To simulate this blue paint BRDF we needed two clear dielectric Torrance-Sparrow layers, and an underlying blue Oren-Nayar surface. The topmost layer has an index of refraction of 1.3 and an average micro-facet slope of 32°. This layer is responsible for the rather strong off-specular reflection. The underlying layer is another Torrance-Sparrow layer with a refractive index of 1.3, and the micro-facets there have an average slope of 8°. To reproduce the retro-reflective properties evident in the measured BRDF, we used an Oren-Nayar surface as the base substrate. The standard deviation of its orientation angle is 20°.

The reason why this second layer does not produce a strong off-specular reflection is, that at grazing angles a part of the energy is absorbed due to total internal reflection, and the fact that the clear varnish we used is – just like a real varnish – not perfectly transparent, but turns successively opaque at very shallow grazing angles.

6.3 Performance Considerations

Since our model consists of a combination of various BRDF models which all have to be evaluated, its computation time mainly depends on the combined evaluation time of the individual BRDFs in the different layers. The combination of the reflectance information from the individual layers does not add significant overhead; even though we did not investigate this, we assume that our technique could be made to work in real time if the individual BRDFs were themselves capable of real time rendering.

In the typical high-quality global illumination setting for which the technique was developed, the evaluation of the reflectance models is an important, but not crucial aspect for the overall performance. For the sake of comparison, we computed the scenes shown in figures 7 and 8 with all layered surfaces being replaced with a normal Torrance-Sparrow surface of equivalent roughness. The inclusion of our model did not significantly increase the execution times of the plain and bi-directional path tracing computations we used to generate the result images in this paper.

7 Conclusion

In this paper we have presented a intuitive and compact multi-layered physically plausible BRDF, which includes absorption as well as total internal reflection between the different layers. However, our approach fails to reproduce wave effects like e.g. iridescence, but is on the other hand far more simple and easier to implement that common algorithms. For those surfaces where we wanted to match a given appearance, our parameters were hand-fitted, which was sufficient to show the capabilities of our approach. Although we needed only a few iterations to find the appropriate values, we will attempt to devise an automatic fitting algorithm for our model in order to increase its usefulness in matching real-world environments.

References

ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic Phong

(a) **(b)** **(c)**

Figure 8: Left: Examples of surface type **a)** applied to crayons and two mugs. Middle: Examples of surface types **d)** (metal foil, the sphere in the background) and **f)** (frosted metal, the two spheres in the foreground). Right: Two copper vases without (left) and with patina (right). The one with patina is an example of surface type **g)**.

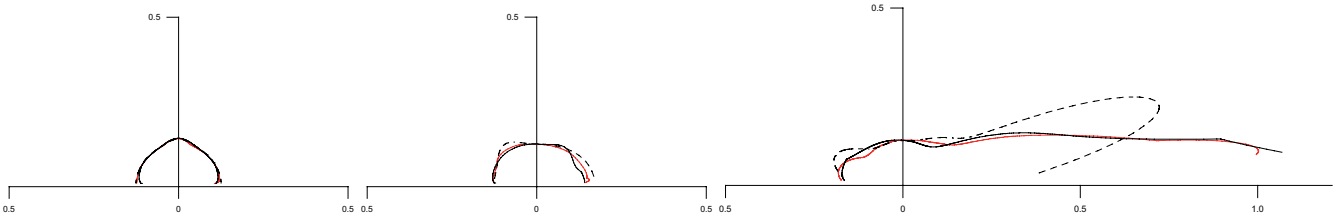

Figure 9: Plots of our reflectance model (solid lines) against the Lafortune model (dashed lines) and the original measured BRDF (red lines) data of blue paint. The incident angle is 0°, 35° and 65° at 550nm. The parameters for the Lafortune model are from [Lafortune et al. 1997].

BRDF model. *Journal of Graphics Tools: JGT 5*, 2, 25–32.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1, 7–24.

Cornell BRDF measurement database. `http://www.graphics.cornell.edu/online/measurements/reflectance/`.

DORSEY, J., AND HANRAHAN, P. 1996. Modeling and rendering of metallic patinas. In *SIGGRAPH*, 387–396.

GRANIER, X., AND HEIDRICH, W. 2003. A simple layered rgb brdf model. *Graphical Models 65*, 4, 171–184.

HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. *Computer Graphics 27*, Annual Conference Series, 165–174.

HIRAYAMA, H., YAMAJI, Y., KANEDA, K., YAMASHITA, H., AND MONDEN, Y. 2000. Rendering iridescent colors appearing on natural objects. In *Pacific Conference on Computer Graphics and Applications*, 15–22.

HIRAYAMA, H., KANEDA, K., YAMASHITA, H., AND MONDEN, Y. 2001. An accurate illumination model for objects coated with multilayer films. *Computers & Graphics 25*, 3, 391–400.

HIRAYAMA, H., KANEDA, K., YAMASHITA, H., YAMAJI, Y., AND MONDEN, Y. 2001. Visualization of optical phenomena caused by multilayer films based on wave optics. *The Visual Computer 17*, 2, 106–120.

ICART, I., AND ARQUÈS, D. 1999. An illumination model for a system of isotropic substrate - isotropic thin film with identical rough boundaries. In *Rendering Techniques*, 261–272.

ICART, I., AND ARQUÈS, D. 2000. A physically-based brdf model for multilayer systems with uncorrelated rough boundaries. In *Rendering Techniques*, 353–364.

KELEMEN, C., AND SZIRMAY-KALOS, L. 2001. A microfacet based coupled specular-matte brdf model with importance sampling. In *Eurographics Short Presentations*, 25–34.

KUBELKA, P., AND MUNK, F. 1931. Ein Beitrag zur Optik der Farbanstriche. *Z. tech. Physik 12*, 593–601.

LAFORTUNE, E. P., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *SIGGRAPH*, 117–126.

NEUMANN, L., AND NEUMANN, A. 1989. Photosimulation: Interreflection with arbitrary reflectance models and illuminations. *Comput. Graph. Forum 8*, 1, 21–34.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of brdf models. In *Proceedings of the Eurographics Symposium on Rendering*, Eurographics Association, 117–226.

OREN, M., AND NAYAR, S. K. 1994. Generalization of lambert's reflectance model. In *SIGGRAPH*, 239–246.

SCHLICK, C. 1993. A Customizable Reflectance Model for Everyday Rendering. In *Fourth Eurographics Workshop on Rendering*, no. Series EG 93 RW, 73–84.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *SIGGRAPH*, 265–272.

WILKIE, A., WEIDLICH, A., LARBOULETTE, C., AND PURGATHOFER, W. 2006. A reflectance model for diffuse fluorescent surfaces. In *GRAPHITE*, ACM, 321–331.