EUROGRAPHICS 2008 / G. Drettakis and R. Scopigno
(Guest Editors)

Volume 27 (2008), Number 2

Photo-realistic Rendering of Metallic Car Paint from Image-Based Measurements

Martin Rump, Gero Müller, Ralf Sarlette, Dirk Koch and Reinhard Klein

Institute for Computer Science II, University of Bonn †

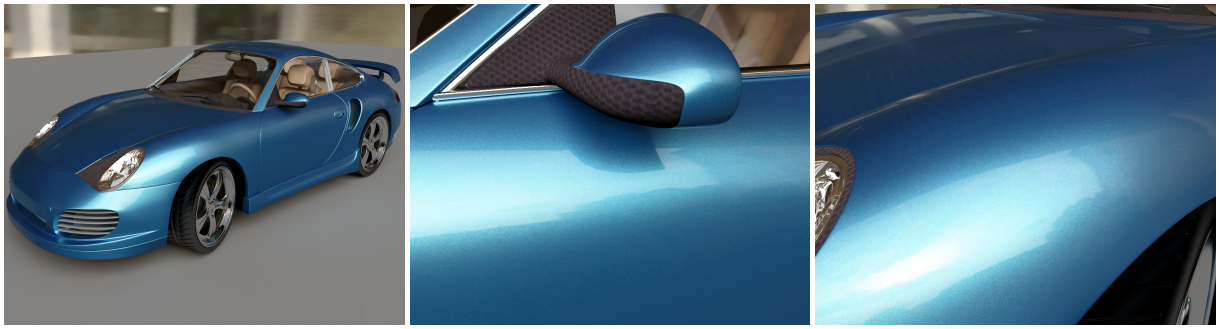

Figure 1: *Based on analytical reflectance modeling and image-based rendering our technique accurately reproduces complex effects of modern car paint like specular reflection, spatially varying glitter with depth impression and color shifts.*

Abstract

State-of-the-art car paint shows not only interesting and subtle angular dependency but also significant spatial variation. Especially in sunlight these variations remain visible even for distances up to a few meters and give the coating a strong impression of depth which cannot be reproduced by a single BRDF model and the kind of procedural noise textures typically used. Instead of explicitly modeling the responsible effect particles we propose to use image-based reflectance measurements of real paint samples and represent their spatial varying part by Bidirectional Texture Functions (BTF). We use classical BRDF models like Cook-Torrance to represent the reflection behavior of the base paint and the highly specular finish and demonstrate how the parameters of these models can be derived from the BTF measurements. For rendering, the image-based spatially varying part is compressed and efficiently synthesized. This paper introduces the first hybrid analytical and image-based representation for car paint and enables the photo-realistic rendering of all significant effects of highly complex coatings.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Picture/Image Generation]: Digitizing and scanning I.3.7 [Three-Dimensional Graphics and Realism]: Color, shading, shadowing, and texture

1. Introduction

In the domain of visual appearance simulation the rendering of metallic car paint is a particular interesting problem. The automotive industry has significant commercial interest in photo-realistic virtual prototyping of their products both during the production process and in the marketing phase where customers want to preview what a specific coating looks like. Moreover, for the rendering community these materials pose a serious challenge because they contain high frequencies both in the angular and in the spatial domain, which makes them difficult to model and measure using standard com-

† e-mail: {rump, gero, sarlette, kochd, rk}@cs.uni-bonn.de

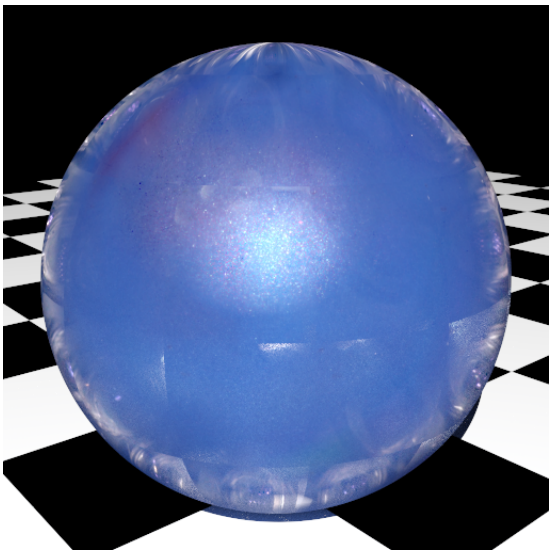

Figure 2: *Rendering metallic car paint using standard BTF measurement and compression techniques. While flake appearance and color shifts are faithfully reproduced, several artifacts like tiling, disturbing reflections at grazing angles and excessively blurred specular reflection destroy the realistic impression.*

puter graphics techniques like BRDF modeling or image-based rendering.

While specialized analytical models for car paint can be derived if a specific physical structure is assumed (e.g., [EKM01]) their practical applicability is limited because of the numerous different coating procedures and structures that are currently used in the industry. Furthermore, the number of parameters of the paint structure grows larger and becomes impractical with the complexity of the paint. Alternatively, general BRDF models like multi-lobe Cook-Torrance [CT81] can be used but this approach suffers from the semantic decoupling of the model parameters from the actual production process which makes it almost impossible to tune the parameters for a given paint. Although the model parameters can be fitted to measurements as in [GCG*05], complex effects like the color shifts in pearlescent paints or the depth impression caused by visible effect particles cannot be modeled accurately by these simple lobe-based models.

If analytical modeling fails (or becomes impractical) computer graphics often relies on image-based rendering. Measuring the reflectance field of a material [DHT*00] allows to reconstruct its appearance for arbitrary viewing and lighting conditions but high-quality samplings of this 8D function are still impractical. Nevertheless, lower dimensional representations like Light Fields [LH96] or Bidirectional Texture Functions [DvGNK97] are already extensively studied and used. Unfortunately, as illustrated in Figure 2 these techniques cannot be straightforwardly applied to highly specular materials like car paint. While the basic appearance of the effect particles and color shifts can be faithfully reproduced, many different types of annoying artifacts like visible repetition due to tiling and reflections of the measurement device on the paint for grazing angles occur. Probably most unpleasant for the experienced observer is the dramatic blurring of the specular lobe due to simple linear interpolation in the undersampled angular domain. Simply increasing the number of images is no solution here because the increase in measurement time and storage space would be infeasible [MMS*05].

To overcome these problems in this paper we separate the appearance of metallic car paint into two parts:

1. The *homogeneous* BRDF part, which describes the reflection behavior of the base and the top layer of the paint and shows high frequencies in the angular domain but stays approximately fixed along the spatial domain. For efficient and accurate capture of the *flip-flop* effect in pearlescent paints (see section 3 for a more detailed explanation) we separate this homogeneous part into other two subparts:
   a. The intensity part, which describes the angular dependent variation in intensity of the paint. This part can be safely modeled using lobe-based BRDF models.
   b. The spectral part, which contains the angular dependent color variations of the paint. These variations show an atypical behavior since lobes are masked for specific directions and wavelengths. Therefore, we build an angular dependent color table from the measurement data.
2. The *spatially varying* BTF part, which is caused by effect particles (typically so-called aluminum flakes) of about 10-50 μm in diameter distributed within the coating, shows high frequencies in the spatial domain but only moderate frequencies in the angular domain. We store this part only after subtracting the homogeneous BRDF part.

We propose to measure these parts simultaneously using a BTF measurement device which captures the spatially varying appearance of a planar material sample with moderate angular resolution. In order to measure also the BRDF part with high angular accuracy we rely on the fact that the incoming and outgoing directions vary across the material sample not only because of a curved shape sample as exploited in image-based BRDF measurement [MWL*99] but simply due to the finite distance of the light source and the image sensor. This provides a dense angular sampling which allows to accurately capture the shape of the material's specular lobe.

Combining the strengths of analytical reflectance modeling and image-based rendering our method is the first which

captures and reproduces (cf. Figure 1) all relevant effects of modern car paints, that are

- surface gloss,
- spatially varying glitter and depth effect,
- color shift (flip-flop),

independent from the actual production process and physical paint structure just from image-based measurements.

2. Previous Work

Modeling of car paint Few papers dealt with analytical models explicitly designed for the paint layer(s).

Ershov et al. [EKM01] build a complete model for the paint layers and their components (binder, pigment particles, flakes, flake coatings). Then they use a technique called "adding or doubling" originally developed to describe planetary atmospheres. This technique includes dividing each paint layer into artificial sublayers which are chosen thin enough, that multiple scattering inside these layers can be neglected. This makes it possible to compute reflection and transmission operators for these layers based on the physical properties of the contained elements. These operators can than be assembled to a reflection operator of the whole paint, which in fact is the BRDF.

To simulate sparkling of flakes in metallic paints Ershov et al. [EKK99] proposed a simple sparkle model. They encounter only flakes that directly reflect light to the observers position. From a normal distribution $P(\beta)$ of the flakes they calculate for a given pixel the probability that a flake has proper alignment for this direct reflection. They scale this probability down to account for flakes deeper in the paint.

Ďurikovič et al. [DM03] explicitly model the geometry for all flakes in the paint. They take flake sizes from real products and let the user define the normal distribution $P(\beta)$ until results look like expected.

Based on these models Ershov et al. [EDKM04] showed how appearance attributes like shade and gloss can be mapped to the physically based paint model parameters to perform reverse engineering of paints.

While already impressive results can be obtained from these techniques a drawback of them is the number of model parameters. For example in the model presented in [EKM01] even one single paint layer can have dozens of parameters depending on the different types of flakes and pigments used and still the depth impression caused by visible effect particles is not realistically conveyed.

These parameter tuning problems have driven the development of methods based on measurements which is also the approach in this paper.

Appearance measurement In computer graphics literature there is a rich branch about measuring appearance for photo-realistic rendering. We will only mention the works with direct relation to the acquisition and rendering of car paint.

The use of measured BRDFs in graphics was introduced by Ward [War92] and the proposed measure-and-fit approach (parameters of an analytic reflectance model are fitted to the measured data) is still widely used. In recent work Ngan et al. [NDM05] qualitatively compared several popular BRDF models. They used a database of about 100 isotropic BRDFs captured by Matusik et al. [MPBM03] which also contains several car paint samples. They concluded that physical based models like the Cook-Torrance model [CT81] (which expresses the specular lobe with respect to the halfway vector and models the Fresnel effect) are better suited for these kind of materials than empirical models like the Ward [War92] or Lafortune [LFTG97] models. Furthermore, they exemplified that more than one lobe is needed for these kind of layered materials.

These findings had some influence on the work of Guenther et al. [GCG*05] who built an BRDF acquisition setup similar to Marschner et al. [MWL*99] in order to obtain sampled car paint BRDFs. They fitted a Cook-Torrance [CT81] BRDF model with 3 lobes to the data which can be directly used for rendering. To simulate the spatially varying appearance of sparkles in metallic paints they added the sparkle simulation from Ershov et al. [EKK99] but stored the normal distribution in a texture to provide frame-by-frame coherence of the sparkles. Their work is probably the most related to the approach presented in this paper. The main difference is that instead of performing BRDF measurements and using a sparkle simulation on top we capture the whole spatially varying appearance using BTF measurements. Furthermore, we explicitly store and model the flip-flop effects of pearlescent paint which are not well represented by lobe-based BRDF models. For the sake of space we will omit a review of BTF research here and refer the reader to the thorough overview about measuring and rendering of spatially varying reflectance by Müller et al. [MMS*05].

Another work on measurement has been done by Sung et al. [SNM*01]. They determine individual flake orientations in the car paint by using a confocal laser scanning microscope and gain a distribution function from this measurements. Additionaly they capture the angular dependent reflectance using a goniometer like setup and build a model for the reflectance based on their measurements.

The rest of the paper is organized as follows: in the next Section we briefly review the basic structure of typical car paint. Then in Section 4 we introduce the BTF measurement procedure we used for appearance capture and explain necessary calibration steps. Sections 5 and 6 introduce the heart of our novel hybrid analytical and image-based model for car paint and how it is rendered effectively. We close the paper with results and conclusions.

Figure 3: *Photograph of a real car door part in bright sunlight. The depth effect and individual flakes remain visible even at several meters distance.*

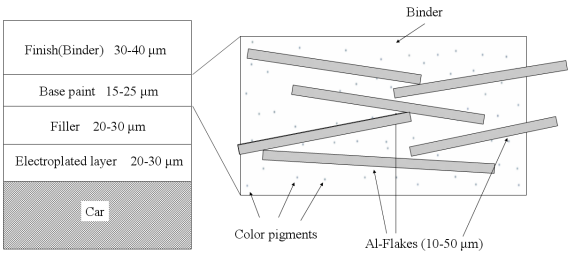

Figure 4: *Typical metallic car paint is composed of different layers. The diameter of the effect particles is usually larger than the thickness of the base paint.*

3. Background

In this section we want to give a short overview over the internal structure of current car paints and explain why the image based approach has advantages compared to an analytical model.

Car paint is a multilayered material. The base layer is the so called *substrate*. At modern cars this is an electroplated layer of tin (as corrosion prevention) covered by a filler made of polished white or light gray powder (to level out the unevenness of the tin) which causes diffuse reflection of incoming light. On top of this filler the main layer of the paint is applied. It consists of binder with color pigments in it. These cause scattering and absorption of incoming light. In the case of metallic paints there are also a large number of disc like metallic particles (called flakes) mixed into this layer. In particular in sunlight these flakes remain prominently visible even for distances in the range of meters (Figure 3). Additionally, these flakes cause a directed diffuse reflection of light which is called *glitter*. An extension to metallic paints are the pearlescent paints (or so called "flip-flop" paints) where the flakes are covered with half transparent layers of mostly metal oxides. These coatings lead to cancelation of certain wavelengths depending on the viewing angle. This results in color shifts dependent on the view and light directions. Many car paints have a second layer (called finish) which is only made of binder. This layer leads to the specular reflection (called *gloss*) of the painted surface. Figure 4 shows the composition of a conceived car paint with scale hints for the components.

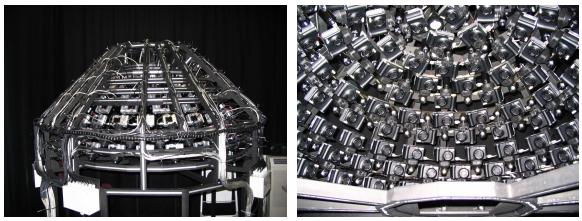

Figure 5: *The BTF measurement device consists of 151 digital consumer cameras mounted on a hemispherical gantry.*

Apart from the precedent components a real car paint production process has many more parameters which influence the final appearance and typically the car paint producers do not make all details of the process publicly available. This makes the design of a predictive and lightweight reflectance model for spatially varying car paint a difficult and almost impossible task. On the other hand using general reflectance models whose parameters have no direct correlation to the parameters of the production process makes manual parameter tweaking impractical since usually more than one lobe is required to represent multi-layer materials [NDM05] which leads to a large set of unintuitive parameters.

4. Measurement

We propose to capture the appearance of spatially varying car paints using BTF measurements. Principally any of the standard devices reviewed in [MMS*05] can be used here. We used a device equipped with 151 digital cameras in order to reduce measurement times and the number of moving parts (see Figure 5). A key feature of the setup is that the built in flashlights are utilized as light sources. Currently we capture 22,801 high-dynamic-range (HDR) images per material which corresponds to typical BTF dataset sizes.

In contrast to pure BTF measurement, where directional lighting and orthographic projection is assumed to be the ideal case, we rely on the fact that in reality image sensor and light source are typically relatively close to the sample. As a result the correct in- and outgoing angles slightly vary across the planar sample. We exploit these variations in order to increase the angular sampling density of the homogeneous BRDF part of the paint. Therefore, we require a very accurate geometrical calibration of the setup. In order to match

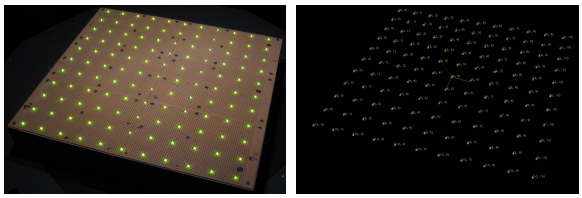

Figure 6: *Left: our calibration object observed under standard room illumination. Right: The same calibration object observed in darkness with overlaid markers over each located LED*

the real paint color as close as possible we also rely on accurate radiometric calibration. In the following subsections we describe the calibration techniques we used to achieve these goals.

4.1. Geometrical Calibration

For the analytical model fitting procedure described in Section 5 we need an as accurate as possible calibration of the intrinsic and extrinsic parameters of each camera. Furthermore, we have to determine the position of each light source. Since the light sources are a fix part of the cameras, we only have to determine position and orientation of the cameras.

The cameras use a small electric motor to move the lens system during zooming and focus operations. Experience has shown that repeated use of this motor can degrade the calibration quality. As a result, when highly accurate calibration information is required all cameras should be recalibrated immediately before or after a sequence of measurements. We use a calibration object specifically designed for that purpose.

Our calibration object exhibits well-known features which can be automatically detected and identified in an image of each camera. We decided on using a planar object with 121 LEDs as features (cf. Figure 6). Compared to typical passive calibration targets the use of LEDs as features when observed in an otherwise dark room makes the detection extremely robust and stable even for grazing viewing angles. Furthermore, we are able to locate the emission center of the LEDs with sub pixel accuracy.

However, a planar calibration object makes it impossible to estimate both extrinsic and intrinsic parameters for every single camera. To solve this, we use a two step calibration procedure. During step I we observe the calibration object without optical zoom and assume the intrinsic parameters of all cameras to be identical. All cameras are of the same type which makes this assumption reasonable. Then we use Zhang's camera calibration algorithm [Zha00] to solve for the distinct extrinsic and the common intrinsic parameters. In step II we consider the extrinsic parameters to be fixed for all cameras and optimize the individual intrinsic parameters of each camera for the desired zoom level using another observation of the same calibration object at the target zoom level. The initial values for the optimization are the intrinsic parameters estimated in step I scaled by zoom dependent factors.

4.2. Dynamic Range

In general, the dynamic range of a reflection on car paint can be as high as the dynamic range of the illumination. Hence the reflection properties have to be measured over a very large dynamic range. We realized this by varying the radiation quantity of the light sources (**F**lash **Q**uantity) and the sensitivity of the CCD sensors (ISO speed). By combining four measurement loops with the settings (FQ2;ISO50), (FQ1;ISO50), (FQ0;ISO50) and (FQ0;ISO400) we achieved a high dynamic range BTF representation of max. $136\,dB$, corresponding to an RGB representation with 23bit per channel.

4.3. Radiometric Calibration

To achieve a consistent color reproduction from all cameras, for each RGB channel of every camera, a response curve was measured taking the unique characteristics of every CCD sensor into account. Slight changes in the characteristics of the sensors are compensated by a white balancing, utilizing a white reflectance standard target.

The spectral distribution of every light source differs slightly. To compensate this, four white reflectance standard targets of different reflectance values from 2.3% to 96.9% are positioned near the paint target during measurement. During post processing an additional white balance for every flash emission is calculated and applied to the affected images.

5. Data Preparation and Modeling

The measurements obtained as described in Section 4 can be used with standard BTF compression and rendering methods (e.g. [SSK03]) but as illustrated in Figure 2 and described in the introduction the results are by far not satisfying.

Therefore, we propose to split the original BTF data into a homogeneous part which can be safely modeled by a BRDF and a spatially varying part which is represented by the measurements minus the BRDF part and can be safely reconstructed using image-based rendering, i.e. interpolation between the samples. In the following we explain this process in detail.

5.1. BRDF Model

As suggested in the study of Ngan et al. [NDM05] we use the well-known Cook-Torrance BRDF model in its multi-lobe version to represent the homogeneous BRDF part of

the car paint. The formula is

$$\rho_X(\omega_i,\omega_o) = \frac{k_d}{\pi} + \sum_{i=1}^{K} \frac{k_{s_i}}{\pi} \frac{F_{F_0,i} D_{m_i} G}{\cos\theta_i \cos\theta_o}. \tag{1}$$

Here ω_i, ω_o are the incoming and outgoing directions. k_d is the diffuse intensity, k_{s_i}, m_i and $F_{0,i}$ are the per-lobe specular coefficient, the distribution exponent and the Fresnel parameter respectively. K is the number of lobes, for the microfacet distribution D we use the Beckmann formulation and the geometric attenuation term G is from [CT81]. For the fresnel term the approximation of Schlick is used. X denotes the whole set of parameters.

While the model is well suited for modeling the glitter and gloss lobes of uniformly colored car paint it has difficulties with the flip-flop effect present in pearlescent paint. The reason is that all lobes basically show a similar behavior, which means they are centered around the ideal reflection direction and vary only by color, the lobe thickness and the Fresnel parameter. This enables for example differently colored glitter and gloss lobes but it fails to express non-diffuse, i.e. view-dependent off-specular color changes. This argument holds also for other microfacet based BRDF models that model mainly the normal distribution of microfacets which is not strongly correlated to the color-shift which results from different physical effects.

For this reason the BRDF is split up into an intensity part described by the Cook-Torrance model and a spectral view and light dependent part. In the case of spectral measurement the intensity would be

$$I = \int_{\lambda=0}^{1} I(\lambda) d\lambda \tag{2}$$

with normalization of the visible wavelength band to $[0,1]$. Having a discrete sampling of this spectrum containing N regular samples, this simplifies to

$$I = \frac{1}{N} \sum_{i=1}^{N} I(i) \tag{3}$$

Since the measurements are in linearized RGB we simplify further and thus are able to divide color and intensity:

$$I = \frac{R+G+B}{3} \tag{4}$$

$$\begin{pmatrix} \acute{R} \\ \acute{G} \\ \acute{B} \end{pmatrix} = \frac{1}{I} \cdot \begin{pmatrix} R \\ G \\ B \end{pmatrix} \tag{5}$$

The BRDF part is fit only at the intensity I and we store the scaled colors in a 4D table χ. Since the color-shifts are low-frequent compared to the gloss reflection a coarse sampling of the colors suffices (we use the angular resolution of the BTF measurement). Compressing the table using for example piecewise polynomial basis functions like B-splines should work well and is left for future work.

The complete BRDF model for the homogeneous part of the car paint is now as follows:

$$\begin{aligned} M_X(\omega_i,\omega_o) \;=\; & \chi(\omega_i,\omega_o) \left(\frac{k_d}{\pi} + \sum_{i=1}^{K-1} \frac{k_{s_i}}{\pi} \frac{F_{F_0,i} D_{m_i} G}{\cos\theta_i \cos\theta_o} \right) \\ & + \frac{k_{s_K}}{\pi} \frac{F_{F_0,K} D_{m_K} G}{\cos\theta_i \cos\theta_o} \end{aligned} \tag{6}$$

The last lobe is used for modelling the gloss of the clearcoat finish and thus is left out of the color modification. Note, that according to equation 3 the whole model works also with spectral measurements, i.e. $N > 3$ color bands. In the following we describe how the model parameters are derived from BTF measurements.

5.1.1. Sample extraction

Since the cameras and light sources are relatively close to the paint sample there is enough angular variation on a single image to capture more than one BRDF data sample from it. Because of the spatial variation resulting from the visible flakes we do not treat any single pixel as a BRDF sample but average across small adaptively chosen patches of the planar probe. To determine the average in- and outgoing directions (ω_i,ω_o) for every patch we require accurate calibration data of the camera and light source positions and orientations (see section 4).

As it turned out during our experiments fitting the whole BRDF model to these samples at once is numerically unstable. In particular the optimizer fails to fit the gloss lobe because there are too few samples of this lobe in the whole sampleset. Therefore, we extract two different sets of samples from the BTF data. The first set represents the diffuse and glitter part of the BRDF while the second set contains data for the gloss lobe.

To extract the diffuse/glitter sampleset from the BTF the following steps are performed:

1. Select BTF images out of the total set of images with a good distribution over the view and the light hemispheres.
2. Based on θ_H a patch size for every image is adaptively chosen: $p := w_{low} + \sqrt{\theta_H/\theta_{H,max}}(w_{high} - w_{low})$. ($\theta_H$ denotes the angle between halfway vector and normal vector, $\theta_{H,max}$ is the greatest accepted value for θ_H of a sample. w_{high} and w_{low} are normally choosen as 70 and 7 pixels for BTF image sizes of 700x700 pixels).
3. Take the average over the patches of the selected size and store it together with the directions obtained from the calibrated light and camera positions. This corresponds to the application of an adaptive box filter to the data.

Averaging over image patches eliminates the noise generated from the flakes. Using an adaptive patch size based on θ_H leads to a better representation of the glitter lobes, since they have the greatest angular variation for small θ_H.

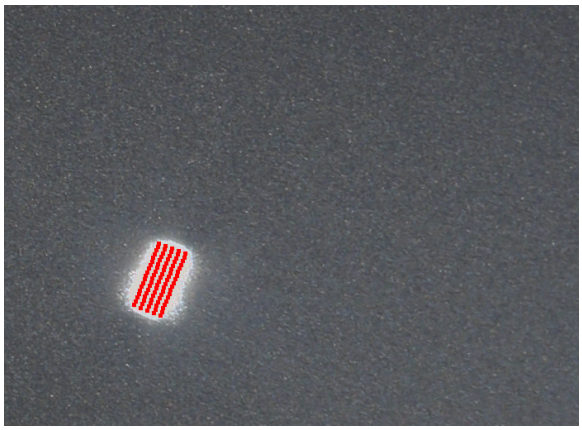

Figure 7: *Accurate calibration is achieved if the virtually mirrored samples of the flash (red dots) coincide with the mirror reflection of the real flash.*

The second sampleset is chosen exclusively from the images which contain the mirror direction and are very accurately calibrated (Figure 7 shows such an image).

We measure the quality of the calibration by comparing the position of the highlight in an image with the reflection position calculated from the calibrated camera and light positions. To find the highlight a list of the n pixels of maximum intensity in the image is built. Then a floodfill with a certain tolerance factor from the pixel of maximum intensity is performed. After the fill an area threshold is checked to test whether the highlight or only a bright flake sparkle was found. If the filled area is too small the next pixel in the list is checked the same way until the highlight is found.

The size of the patches for the second sampleset is chosen as small as possible in order to preserve as much information about the shape of the gloss lobe as possible and big enough to filter out the noise from the flakes. From our experience a patch size of 4x4 pixels (compared to an image size of about 700x700 pixels for the BTF measurement) turned out to be a good compromise.

5.1.2. Fitting process

Having obtained the two sets of samples the actual fitting process can start. We use a Quasi-Newton SR-1 method with the following error metric:

$$E(X) = \sum_{i=1}^{N} \left| log \frac{\Psi(\omega_i, \omega_o) + \varepsilon}{\frac{1}{L} \sum_{l=1}^{L} M_X(\omega_{i_l}, \omega_{o_l}) + \varepsilon} \right| + P(X) \quad (7)$$

Where P is defined as:

$$P(X) := \begin{cases} 0, \text{ if X physically plausible} \\ S, \text{ else} \end{cases} \quad (8)$$

having:

M	Cook-Torrance model
X	parameter set for the model
Ψ	data samples as explained in section 5.1.1
L	number of point light sources to sample the flash
N	total number of samples
P	plausibility function for the parameter set X
S	a big constant
ε	a small constant to prevent division by zero errors which are possible for certain parameters X

As explained in the last section it is advantageous that the cameras are near to the sample. But since the light source (flash) has a certain extent we cannot model the flash by a single point light source. Instead we approximate the shape of the flash by distributing a number of point light sources (e.g. L=90) over the the flash window of the camera.

The fitting process has three steps:

1. Using only the diffuse/glitter sampleset we fit the diffuse intensity and 1-3 glitter lobes. We start from several different initial points to avoid running into a local minimum of the objective function.
2. Using only the gloss sampleset we fit only the parameters of one single lobe for the gloss part while keeping fixed the parameters for the diffuse/glitter part computed in the previous step. Again we start from several initial sets of parameters and store only the parameters yielding the best fit.
3. We iterate step 1 and 2 a few times. The only difference is that the initial parameters used are the fit results from the previous iteration. This helps to reduce errors resulting from the separate fit of the two parts.

With our implementation the whole extraction and fitting process takes about 30 minutes using a machine with an Intel E6600 processor (Core 2 Duo) and 1 GB of RAM. Table 1 shows the parameter sets for the three measured paints and Figure 8 shows a comparison of the fitted BRDF model with the input data for the green-blue flip-flop paint.

5.2. Spatially Varying Model

After the parameters of the Cook-Torrance model are computed by the fit procedure we subtract the model from the captured images. The directions ω_i, ω_o are calculated for every pixel x from the calibration data and the model is evaluated for that directions. Then a simple subtraction in RGB space is done. The resulting images then only contain the flake specific data. Figure 9 shows a small patch of such a difference image.

We call the resulting resampled flake BTF simply $\overline{BTF}$ and end up with the following hybrid analytical and image-based model for car paint:

$$\mathbf{M}_X(x, \omega_i, \omega_o) = \overline{BTF}(x, \omega_i, \omega_o) + M_X(\omega_i, \omega_o) \quad (9)$$

	green-blue flip-flop	silver metallic	specular blue
k_d	0.0201372	0.0717947	0.038105
k_{s_0}	0.140216	0.167496	0.0435869
m_0	0.851337	1.29319	1.04887
F_{0_0}	0.115703	0.442487	0.120618
k_{s_1}	0.0993458	0.226424	0.0834049
m_1	0.248163	0.453529	0.215455
F_{0_1}	0.172867	0.558767	0.108905
k_{s_2}	0.145249	0.290604	0.030003
m_2	0.142194	0.227447	0.00282852
F_{0_2}	0.285121	0.807962	0.0351793
k_{s_3}	0.15	0.2	-
m_3	0.003	0.002	-
F_{0_3}	0.0851148	0.0942766	-

Table 1: *BRDF parameters for the paints*

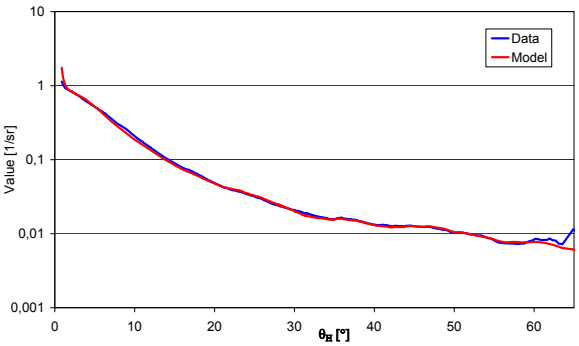

Figure 8: *Comparison of BRDF values between data (blue) and model (red) for the green-blue flip-flop paint using 4 lobes*

6. Synthesis and Rendering

The model from equation 9 in its raw form is not suitable for direct rendering because simple tiling of the otherwise randomly distributed flakes would be clearly visible and disturbing. Furthermore, the uncompressed size of the BTF part for a 128x128 pixel patch with the full set of 151x151 directions would be more than 2 GB which is far too large for effective rendering.

The second problem can easily be alleviated by employing BTF compression algorithms. Currently, we compress the flake data $\overline{BTF}$ using the per-view factorization method from [SSK03] with a number of 20-25 PCA components which reduces the storage requirements to about 500 MB (including MIP-map levels).

To deal with the tiling problem we propose simple copying of random patches (of size about 16x16 texels). This synthesis can be done on the fly during the rendering process and performs surprisingly well. Figure 10 shows that our approach keeps the overall impression of the paint. To provide

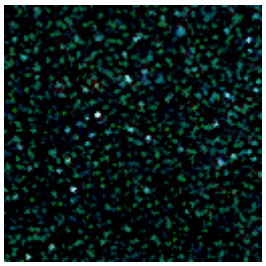

Figure 9: *Image of the spatially varying flake data $\overline{BTF}$ after subtracting the homogeneous BRDF part.*

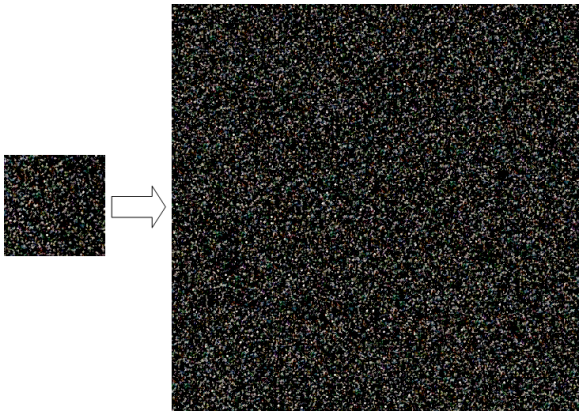

Figure 10: *Example of synthesis by copying random patches*

frame-to-frame coherence the mapping between source and destination patches can be stored an reused. Usually we keep only 128x128 pixels of the BTF data which correspond to an extent of 1,2x1,2 cm^2. Our experience has shown that such a small patch of car paint holds already enough variance to mask any tiling artifacts.

In order to avoid aliasing artifacts for farther viewing distances we prefilter the BTF part in the spatial domain using a simple box filter and use MIP-mapping during rendering.

For rendering within a global illumination framework we use importance sampling of the BRDF part and the environment map to generate appropriate outgoing directions. Note that for metallic paint it is important to sample the whole BRDF and not only the gloss lobe because the paint layer typically collects significant amounts of light also from off-specular directions. We use the method from [LRR05] to sample both the BRDF and the environment map and use a balance heuristic for the combination of both.

In order to convert the resulting images from HDR to displayable or printable LDR we use the measured response curve of one camera from the measurement setup instead of a tonemapping operator.

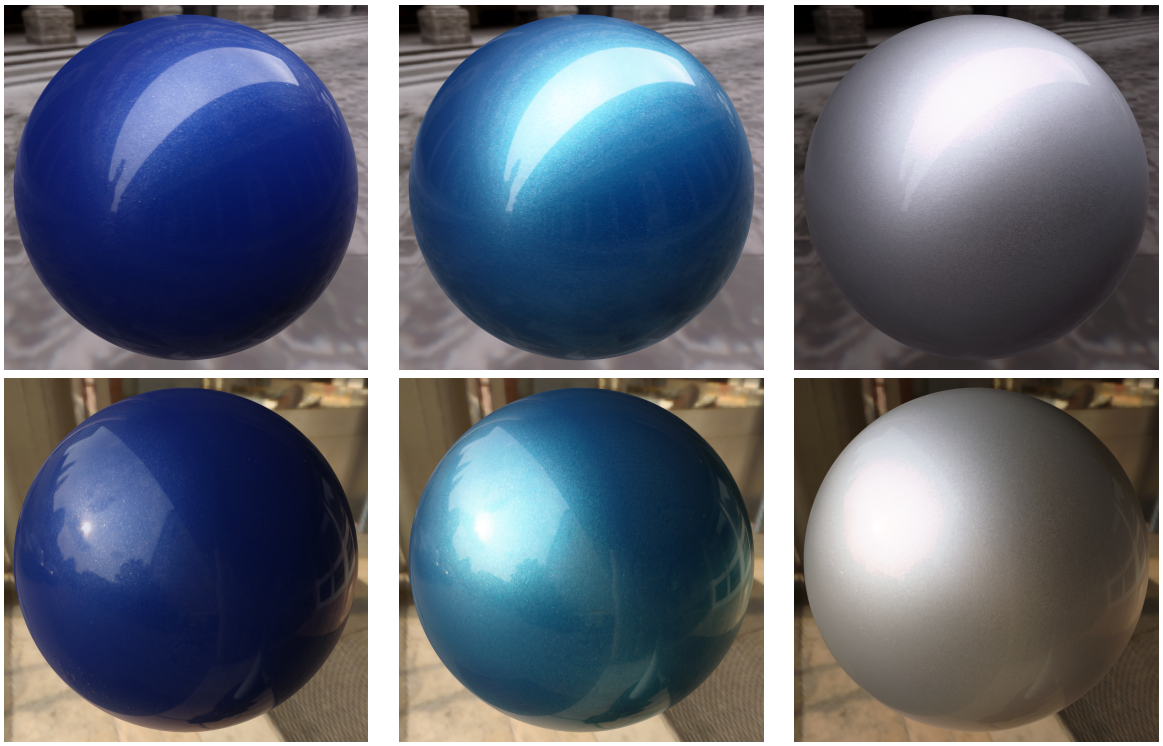

Figure 11: *The materials used in this paper lit by different environment maps choosing an appropriate exposure per image. Left: specular blue, middle: green-blue flip-flop, right: silver metallic. The flake density in the paint increases from left to right.*

Figure 12: *The Porsche model covered with green-blue flip-flop paint rendered with the main entrance environment map.*

Figure 13: *The main entrance environment map.*

7. Results

We applied the data preparation and fitting process described in Section 5 to 3 different car paints which are depicted in Figure 11. We used two different environment-maps here, the first is the well known Uffizi environment from Paul Debevec and the second was captured by ourselves and is shown in Figure 13. It is interesting, because it contains direct sunlight. The synthesis and rendering stage described in Section 6 was implemented as a MentalRay-shader.

The specular blue paint has an almost mirror-like finish and a moderate amount of visible flakes which can be seen in Figure 11. The green-blue flip-flop paint shows the probably most interesting effects. It has specular reflection, a serious amount of visible flakes and also a slight green-blue color shift resulting from flip-flop effect particles. The silver

metallic paint has a paint layer with a very large amount of flakes which leads to very prominent glitter reflection.

8. Conclusions and Future Work

We presented a novel measurement and rendering framework especially designed for metallic car paint and based on BTF measurements. The measurements capture important effects like depth impression and color-shifts which are preserved in our hybrid analytical and image-based model. The highly specular parts of the material are represented with a BRDF model whose parameters are fitted to the data using a novel BTF resampling procedure. Rendering is performed using BRDF and BTF rendering techniques.

In future work we plan to investigate more sophisticated BTF compression techniques. Concerning predictive rendering under complex illumination, spectral measurement and rendering is an important issue worth of additional research.

References

[CT81] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. In *SIGGRAPH '81: Proceedings of the 8th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1981), ACM Press, pp. 307–316.

[DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. *Proceedings of ACM SIGGRAPH 2000* (2000).

[DM03] DURIKOVIČ R., MARTENS W. L.: Simulation of sparkling and depth effect in paints. In *SCCG '03: Proceedings of the 19th spring conference on Computer graphics* (New York, NY, USA, 2003), ACM Press, pp. 193–198.

[DvGNK97] DANA K. J., VAN GINNEKEN B., NAYAR S. K., KOENDERINK J. J.: Reflectance and texture of real-world surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition* (1997), pp. 151–157.

[EDKM04] ERSHOV S., DURIKOVIC R., KOLCHIN K., MYSZKOWSKI K.: Reverse engineering approach to appearance-based design of metallic and pearlescent paints. *The Visual Computer 20*, 8-9 (October 2004), 587–600.

[EKK99] ERSHOV S., KHODULEV A., KOLCHIN K.: Simulation of sparkles in metallic paints. In *Proceedings of Graphicon '99* (1999), The Eurographics Association and Blackwell Publishers, pp. 121–128.

[EKM01] ERSHOV S., KOLCHIN K., MYSZKOWSKI K.: Rendering pearlescent appearance based on paint-composition modelling. In *EG 2001 Proceedings*, Chalmers A., Rhyne T.-M., (Eds.), vol. 20(3). Blackwell Publishing, 2001, pp. 227–238.

[GCG*05] GÜNTHER J., CHEN T., GOESELE M., WALD I., SEIDEL H.-P.: Efficient acquisition and realistic rendering of car paint. In *Vision, Modeling, and Visualization 2005 (VMV'05)* (Erlangen, Germany, November 2005), Greiner G., Hornegger J., Niemann H., Stamminger M., (Eds.), Aka, pp. 487–494.

[LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proceedings of SIGGRAPH* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 117–126.

[LH96] LEVOY M., HANRAHAN P.: Light field rendering. *Computer Graphics 30*, Annual Conference Series (1996), 31–42.

[LRR05] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHI R.: Adaptive numerical cumulative distribution functions for efficient importance sampling. In *Eurographics Symposium on Rendering* (June 2005).

[MMS*05] MÜLLER G., MESETH J., SATTLER M., SARLETTE R., KLEIN R.: Acquisition, synthesis and rendering of bidirectional texture functions. *Computer Graphics Forum 24*, 1 (March 2005), 83–109.

[MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: Efficient isotropic BRDF measurement. In *Proceedings of the EGRW '03* (2003), pp. 241–247.

[MWL*99] MARSCHNER S. R., WESTIN S. H., LAFORTUNE E. P. F., TORRANCE K. E., GREENBERG D. P.: Image-based BRDF measurement including human skin. In *Proceedings of 10th Eurographics Workshop on Rendering* (1999), pp. 139–152.

[NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of brdf models. In *Proceedings of the Eurographics Symposium on Rendering* (2005), Eurographics Association, pp. 117–226.

[SNM*01] SUNG L. P., NADAL M. E., MCKNIGHT M. E., MARX E., DUTRUC-ROSSET R., LAURENTI B.: Effect of aluminum flake orientation on coating appearance. *Proceedings of the 79th Federation of Societies for Coatings Technology (FSCT ICE) Meeting* (nov 2001), 1–15.

[SSK03] SATTLER M., SARLETTE R., KLEIN R.: Efficient and realistic visualization of cloth. *Proceedings of the Eurographics Symposium on Rendering 2003* (2003).

[War92] WARD G. J.: Measuring and modeling anisotropic reflection. In *Proceedings of SIGGRAPH* (1992), ACM Press, pp. 265–272.

[Zha00] ZHANG Z.: A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence 22*, 11 (2000), 1330–1334.