

Efficient Acquisition and Realistic Rendering of Car Paint

Johannes Günther, Tongbo Chen, Michael Goesele, Ingo Wald, and Hans-Peter Seidel
MPI Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany
{guenther, tongbo, goesele, wald, hpseidel}@mpi-inf.mpg.de



Figure 1: Different views of a “VW New Beetle” car rendered with a measured and fitted “Polaris Silber” car paint, lighted from Paul Debevec’s HDR environment map of the Uffizi Gallery in Florence [5].

Abstract

The outside appearance of cars is mostly defined through only two distinct materials – glass and car paint. While glass can rather easily be simulated by the simple physical laws of reflection and refraction, modelling car paint is more challenging.

In this paper we present a framework for the efficient acquisition and realistic rendering of real-world car paint. This is achieved by building an easy-to-reproduce measuring setup, fitting the measured data to a general BRDF model for car paint, adding a component for simulating the sparkling effect of metallic paints, and rendering using a specially designed shader in a realtime ray tracer.

1 Introduction

One of the most important applications of computer graphics in industry is “virtual prototyping”, i.e., the process of designing an industrial product (e.g., a car) completely digital, and basing all design decisions solely on computer-generated images of that virtual car. Obviously, this is only feasible if the rendered images are sufficiently close to reality. In particular for the automotive industry, this requires to simulate two distinct kinds of material, glass and car paint. Glass – which is important for both the windshields and for the car’s lights – can be easily simulated using the well-known physical laws of reflection and refraction. However, in particular metallic car paint is more challenging, as it has many different components, including a diffuse component, glossiness, almost perfectly specular reflections, highlight, glitter, etc.

In particular for interactive applications, car paint is usually represented using simple reflection models such as the Phong [20] model, with the model’s parameters tweaked by the user to look “nice”. This however is both ad-hoc and laborious, and often fails to produce a realistic appearance; in practice, rendered cars often appear to be made of plastic.

We address this problem by measuring samples of real car paint to reproduce the look as close as possible. We then investigate advanced reflection models for car paint, and determine their parameters for real-world (pearlescent) car paints by fitting the measured data. The fitted car paint model is then realized using a specially designed car paint shader that is integrated into a realtime ray tracer. This shader – apart from simulating glossy and specular components – also reproduces the glitter effect caused by the flakes in the metallic car paint. Finally, as a realistic appearance of car paint is also strongly affected by the way it is lighted, we light the car using captured HDR environmental lighting. Thus, our framework covers the whole span from efficient acquisition of real-world car paint to its realistic (interactive) rendering.

Contributions: In this paper, we present

- An image-based BRDF acquisition setup that is designed to be easily reproducible.
- An analysis of the acquired data, derivation of a suitable car paint BRDF model, and determination of its parameters for various car paints.
- A framework for realistically rendering the acquired data including diffuse, specular, and glossy components as well as sparkles.

2 Previous Work

In this section, we will briefly describe the general ideas and principles that our work is based on.

Marschner et al. [16] proposed an efficient method to measure the isotropic BRDF of materials. A homogeneous sphere sample of the target material is captured with a digital camera in different lighting conditions. The curved sample surface allows to acquire many BRDF samples at once and yields a short acquisition time. Matusik et al. [17] built an automatic acquisition system and captured a large database of various materials. They analyzed the gathered data and derived a data-driven BRDF model. Ngan et al. [19] analyzed how various analytic BRDF models perform when be used to fit to real measured data from [17]. As one result they found that the physically based BRDF model of Cook and Torrance [4] performed very well, especially with metal-like materials.

Ershov et al. [7, 8] presented a complex physically based multi-layer paint model, particularly concentrating on modeling the reflection behavior of car paints. The simulation of the sparkling effect present in metallic paints was also addressed by Ershov et al. [9] as well as by Ďurinkovič et al. [24].

Capturing the illumination conditions of places of the real world was demonstrated by Debevec et al. [5]. Algorithms for the efficient sampling of the resulting environment maps in order to render virtual scenes under the captured illumination were proposed by Agarwal et al. [1] as well as by Kollig and Keller [14].

Cabral et al. [3] proposed an image based algorithm to quickly render complex BRDFs, in particular car paints. However the costly preprocessing to generate radiance environment maps does not allow for changing the lighting, scene composition, or BRDF interactively.

3 BRDF Acquisition

Traditionally, bidirectional distribution functions (BRDFs) are measured by using specially designed devices such as, e.g., gonireflectometers in a restricted laboratory environment. In order to accelerate the acquisition process, image-based measurement techniques use digital cameras as measurement devices, which – due to the large number of pixels in an image – allows for acquiring




Figure 2: Top: A photograph of our BRDF measurement system. Bottom left: Our measurement target, a sphere painted with car paint. Bottom right: A sample HDR image which is gamma corrected and scaled for illustration purpose.

many BRDF samples in a single image. Inspired by Marschner et al. [16] and Matusik et al. [17], we have built our own high-speed BRDF measurement system (see Figure 2), which includes an accurate turn table, a high-quality digital camera, a point-like light source, and a painted sphere as acquisition target (see also Figure 9). The sphere has been covered with car paint by a professional painting studio, and is mounted on the center of the turn table, and is attached to the turn table is an arm with the light source fixed to it; the camera is mounted statically. The whole system is placed in a dark photo studio, whose walls and ceiling are covered with black felt; the floor is covered with black needle fleece carpet [10].

To ensure fast and accurate acquisition, we choose the camera and light source carefully. We use a 5 Watts Luxeon white LED with an additional optics as stable point light source. The focusing optics mounted on the LED shortens the acquisition time significantly and ensures uniform illumination of the sample sphere. For capturing images, we choose a high quality scientific CCD digital camera, a Jenoptik ProgRes C14. This camera captures high quality images with low noise even at long exposure times.

Before the acquisition, we do both geometric and radiometric calibration for the system. In geometric calibration, we first recover the intrinsic param-

eters of the camera by using MatLab’s Camera Calibration Toolbox [2]. The radius of the painted sphere is measured using a caliper. We then calculate the center of the sphere with the known parameters, and setup the geometric relationship of camera, light source, and sphere. For radiometric calibration, we recover the response curve of the camera, which describes the relationship between pixel values and luminance (irradiance) [22]. The benefit of this method is to do optimization over all pixels in all images with an adequate weighting scheme.

During acquisition, the light source moves in increments of 1 degree from the point exactly in front of the camera to the position exactly opposite the camera. For each position of the light, we take 16 images with exposure times ranging from 0.25ms to 328ms and combine them to generate one HDR image, which records the BRDF data for this light position, and in which each pixel on the sphere is a separate BRDF measurement. To avoid considering reflections of arm and turn table, only the upper half-sphere is taken into account. The ratio of the high dynamic range radiance to the irradiance is treated as the BRDF value. By using basic geometry, we can get the position on the sphere, the incoming light ray, the viewing direction and the normal of each pixel. Finally we record the measured BRDF samples in a compact BRDF table using an efficient halfway vector based, nonlinear parameterization (see Section 4.1).

4 Data Representation and Modelling

4.1 Parameterization

For highly specular car paint, the specular peaks are difficult to represent using the traditional coordinate system. We use a halfway vector (the vector halfway between the directions of incidence and reflection) based parameterization [23], which can represent BRDF compactly and allows us to sample the highlight more densely than the off-highlight BRDF. Figure 3 shows these two parameterization schemes. We discretize θ_h and θ_d into 90 bins and ϕ_d into 180 bins. The empty bins (holes) of the BRDF table are filled with weighted sum of the neighbor BRDF samples. The size of the uncompressed BRDF table with float precision is about 18 Mbytes.

Figure 4 shows results of different parameterization schemes. We found the halfway vector based

parameterization with nonlinear sampling of θ_h performs best, in particular for representing the specular highlights.




Figure 3: Different ways to parameterize an isotropic BRDF. Left: With respect to the normal $f_r(\theta_i, \theta_o, \phi_d)$. Right: With respect to the halfway vector $f_r(\theta_h, \theta_d, \phi_d)$.




Figure 4: Comparison of different BRDF parameterizations. Top Row: Without interpolation and hole-filling, empty bins marked with black. Bottom Row: With interpolation and hole-filling. Left: Traditional BRDF parameterization. Center: Linear halfway vector based parameterization. Right: Halfway vector based parameterization with nonlinear sampling of θ_h .

4.2 Fitting to Cook-Torrance

The Cook-Torrance model [4] is physically derived and has shown to perform well with many materials [19]. In its multi-lobe form, the Cook-Torrance BRDF can be expressed as:

$$f_r(x, \omega_i, \omega_o) = \frac{kd}{\pi} + \sum_{i=1}^N \frac{ks_i F_{r_i} D_{m_i} G}{\pi (\mathbf{n} \cdot \omega_i) (\mathbf{n} \cdot \omega_o)}, \quad (1)$$

where N is the number of specular lobes, kd and ks are the diffuse and specular parameters, respectively, and m and r are the parameters for the distribution and Fresnel terms, respectively. We treat

kd and ks as color channel dependent RGB floats, while m and r are scalars.

For fitting, we use the constrained Levenberg-Marquardt method [21] to solve the optimization problem. The energy function used during the optimization is

$$E(X) = \sum_{\Theta} |\log(\mathcal{C}(X, \Theta)/\mathcal{D}(\Theta))| + \Psi(X),$$

where \mathcal{C} is the fitted BRDF value, \mathcal{D} is the measured BRDF, Θ represents the angular parameters of the BRDF sample, Ψ is the penalty term, and $X = \{kd, ks_1, m_1, r_1, \dots, ks_N, m_N, r_N\}$ represents the parameters to be fitted. The penalty term $\Psi(X)$ is defined as

$$\Psi(X) = \begin{cases} 0, & X \text{ is plausible} \\ K, & X \text{ is implausible} \end{cases} \quad (3)$$

where “plausibility” is based on physical plausibility of the parameters of the Cook-Torrance model (i.e., both m and r in X should be in $(0, 1)$). The constant K is chosen heuristically.

Another issue in stable fitting is the difficulty of effectively approximating the specular peak. The specular highlight tends to be very sensitive to numerical and measurement errors, as such a prominent and narrow feature of the function is hard to approximate. McCool et al. [18] use a separate specular texture outside of his approximation to improve the appearance and we also model the specular peak separately using a third specular lobe.

Figure 5 shows the fitting results with a two-lobe Cook-Torrance model for the “Polaris Silber” BRDF data. In order to fit the Fresnel parameters effectively, we sample the BRDF table with respect to the angle between halfway vector and surface normal. The fitted parameters for all our measured car paints can be found in table 1.

5 Realistic Rendering

Having obtained the data of car paint, either as BRDF table or as a parameter set of an analytic BRDF model, we now describe how to use this data to render realistically looking images efficiently.

For this purpose we decided to use ray tracing which recently became fast enough to provide interactive frame rates [27, 26]. Ray tracing has the additional advantage to be flexible, and accurate when calculating reflections.

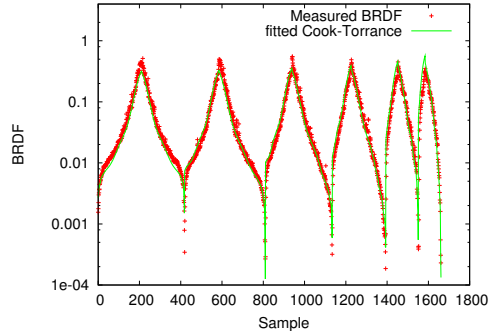


Figure 5: Fitting example of measured “Polaris Silber” BRDF data with a two-lobe Cook-Torrance model.

5.1 Complex Illumination

5.1.1 HDR Environment Maps

In order to have a realistic lighting environment for the car model we use captured environment maps. To obtain the illumination of one place we follow the approach of Debevec et al. [5] and use these high dynamic range environment maps to light the scene. Consequently we have also to deal with the high dynamic range of the rendered images. Tonemapping is the task to represent this high dynamic range on low dynamic range displays appropriately. Many different algorithms were proposed in the past (see e.g. [6] for an overview), unfortunately most of them are not suited for realtime rendering or are problematic with animations.

To circumvent these problems we choose a simple non-linear compression of the dynamic range, which works well for our case. Alternatively GPU implementations of tonemapping operators could have been used as well.

5.1.2 Sampling Strategies

For the actual rendering we have to evaluate the following integral [11] to shade a surface point x with normal \mathbf{n}

$$L_o(x, \omega_o) = L_e(x, \omega_o) + \int_{\Omega} f_r(x, \omega_i, \omega_o) L_i(x, -\omega_i)(\mathbf{n} \cdot \omega_i) d\omega_i,$$

where $L(x, \omega)$ is the radiance at point x in direction ω , and f_r is the BRDF.

parameter	Ford F8	Polaris Silver	Opel Titan	BMW 339
kd	.0012, .0015, .0018	.055, .063, .071	.011, .013, .015	.012, .015, .016
ks_1	.0049, .0076, .012	.065, .082, .088	.057, .066, .078	.062, .076, .08
m_1	.32	.38	.38	.39
r_1	.15	1	.85	.92
ks_2	.01, .013, .018	.11, .13, .13	.1, .12, .13	.11, .12, .12
m_2	.11	.17	.17	.17
r_2	.087	.92	.86	.87
ks_3	.007, .0065, .0077	.008, .013, .015	.0095, .014, .016	.0083, .015, .016
m_3	.013	.013	.014	.013
r_3	.9	.9	.9	.9
	2K Acrylack	White	Blue	Blue Matte
kd	.42, .32, .1	.61, .63, .55	.0079, .023, .1	.0099, .036, .12
ks_1	0, 0, 0	2.6e-06, .00031, 3.1e-08	.0011, .0015, .0019	.0032, .0045, .0059
m_1	.88	1	.15	.16
r_1	1	.049	1	1
ks_2	.028, .026, .006	.013, .011, .0083	.025, .03, .043	.18, .23, .28
m_2	.8	.15	.043	.075
r_2	.9	.45	.094	.046
ks_3	.017, .0075, .0041	.049, .042, .037	.059, .074, .082	.04, .049, .051
m_3	.015	.015	.02	.034
r_3	.17	.17	.17	.17

Table 1: Fitted parameters of the multi-lobe Cook-Torrance BRDF model to the measured BRDF data for several car paints. Rendering results using these parameters can be found in Figure 10.

In computer graphics many ways have been proposed to calculate this integral. A straightforward way would be to use Monte-Carlo integration in the form of path tracing [15, 25]. However, if we have only little information about the BRDF f_r we can only do importance sampling of the cosine term resulting in either very noisy images or prohibitive many samples and thus huge rendering time.

Fortunately we can use more efficient integration methods. Since we use environment maps to illuminate the scene direct illumination will dominate. Neglecting diffuse indirect illumination allows for importance sampling of the incident light $L_i(x, -\omega_i)$. Therefore, the environment map is discretized into l directional lights $\{L_d, \omega_d\}_{d=1}^l$ by methods such as [1, 14]. The shading of point x is thus approximated by

$$L_o(x, \omega_o) \approx \sum_{d=1}^l f_r(x, \omega_d, \omega_o) I_d(\mathbf{n} \cdot \omega_d). \quad (4)$$

This procedure has the advantage that we now need fewer illumination samples (the important ones) for shading reducing computational costs. And because these directional light samples will be used for shading every point the noise due to the variance of the Monte-Carlo estimate will be suppressed at the same time.

However, there are also some drawbacks. If the BRDF f_r is highly specular the discretization of the

environment map will cause visible artifacts, i.e. a glossy highlight will be split into several small highlights corresponding to the individual directional lights. In this case it is advantageous to do importance sampling of the BRDF instead of the illumination. Importance sampling means to generate s random directions $\{\omega_j\}_{j=1}^s$ according to a probability distribution function pdf that ideally is proportional to f_r , and to gather the radiance coming from these directions using ray tracing. Thus we get

$$L_o(x, \omega_o) \approx \sum_{j=1}^s \frac{f_r(x, \omega_j, \omega_o)}{pdf(\omega_j)} L(rt(x, \omega_j), -\omega_j), \quad (5)$$

where $rt(x, \omega)$ is the ray tracing function returning the closest hit-point to x in direction ω . As this sampling strategy is only used for highly specular BRDFs only few samples are necessary without introducing too much noise.

5.1.3 Application to Car Paint BRDFs

In Section 4.2 we already showed that the BRDF of typical car paints is a mix of a diffuse part, a broad gloss, and a spiky specular part, hence we can neither apply (4) nor (5) directly. However by splitting the BRDF into a moderate glossy part (including the diffuse part) and a highly specular part we can use different sampling strategies for each of the individual parts.

This splitting can easily be done with the fitted multi-lobe Cook-Torrance BRDF after identifying the lobe modelling the highly specular peak caused by the clear coat.

In order to find an appropriate *pdf* to efficiently sample the highly specular lobe according to (5) we approximate the distribution term in Cook-Torrance model by a cosine lobe, which was also proposed by Kelemen [12]. Thus we approximate

$$D(\theta) = \frac{1}{4m^2 \cos^4 \theta} e^{-\frac{\tan^2 \theta}{m^2}} \approx \frac{n+1}{2\pi} \cos^n \theta,$$

$$\text{where } n = \frac{2\pi}{4m^2} - 1.$$

Given two uniformly distributed random numbers r_1 and r_2 we can now obtain a random halfway vector $\mathbf{h}(\theta, \phi)$, $\theta = \arccos r_1^{\frac{1}{n+1}}$ and $\phi = 2\pi r_2$, with distribution $D(\mathbf{h} \cdot \mathbf{n})$. After mirroring the outgoing direction ω_o onto \mathbf{h} we get one incoming sample direction ω_j . Because of this transformation the *pdf* needs to be divided by $4(\mathbf{h} \cdot \omega_o)$.

For the tabulated BRDF the splitting into a moderate glossy and a highly specular part is not that straightforward. One possibility is to clamp the values of the BRDF table to the maximum value of the glitter lobe to end up with only moderate glossy part, effectively cutting the specular peak. To account for the specular peak itself a reflection ray into the specular direction is sent to gather the incident light of this direction. Alternatively one can try to use the fitting parameters to subtract the highly specular lobe from the BRDF data to get the moderate glossy part of the BRDF. The specular peak can then be sampled as described above.

5.2 Simulation of Sparkling

The sparkle effect is an interesting feature of metallic car paints and one reason for the popularity of these paints. This effect materializes as tiny bright spots when viewed from close distance (see Figure 2b and c) – especially in sunlight. It is caused by small flakes embedded in the paint layer. The randomly oriented flakes act as small mirrors reflecting the light directly to the observer.

When the virtual car is viewed from distance the individual sparkles are not visible and their (integrated) contribution is already contained in the measured BRDF. Thus the normal evaluation of the BRDF is totally sufficient. However, at close proximity it is desirable to mimic the sparkling effect.

To simulate this effect Ershov et al. [9, 7] proposed a detailed physically based model of the flakes and the actual rendering was driven by derived statistically properties. Đurinkovič et al. [24] took a different approach and explicitly modeled the flakes with geometry generated regarding statistical properties of the flakes such as density and distribution. We will use ideas from both approaches.

As we want to animate the virtual car or inspect the car paint interactively it is crucial to maintain the spatial and temporal coherence of the sparkling pattern. Therefor a purely statistical approach drops out – the flakes need to be modeled explicitly to some extend. On the other hand it is prohibitive to create geometry for millions of flakes for the entire car due to the huge storage and computational costs.

Our solution meeting these requirements is the use of a procedural texture map holding the normals of the flakes: The (u, v) texture coordinates of the to be shaded point are used to initialize a pseudo random number generator.

From the fitted parameters of the measured BRDF to the micro-facet based Cook-Torrance BRDF model we can derive a distribution function of the flake normals. In order to obtain a sampling scheme for the (pseudo random) flake normal we approximate the Beckmann distribution of the Cook-Torrance model as follows [12]

$$D(\theta) = \frac{1}{m^2 \cos^4 \theta} e^{-\frac{\tan^2 \theta}{m^2}} \approx \frac{1}{m^2 \pi \cos^3 \theta} e^{-\frac{\tan^2 \theta}{m^2}}.$$

Using two (pseudo) random numbers r_1 and r_2 the flake normal $\mathbf{n}_f(\theta, \phi)$ is calculated by $\theta = \arctan\left(m\sqrt{-\log(1-r_1)}\right)$ and $\phi = 2\pi r_2$. Having the normal of the flake it is straightforward to illuminate the highly glossy flake.

Aliasing artifacts may occur because the flakes have a finite size but are very small. These artifacts can be suppressed by oversampling or a distance-dependent blending function. Figure 6 shows the results of our sparkle simulation, demonstrating the maintained frame-to-frame coherence as well as the influence of the viewing distance.

5.3 Optimizing Quality and Speed

So far we have mainly concentrated on achieving the best possible rendering quality. However, we obviously want to apply our improved car paint model in an interactive setting as well. Fortunately, there are several parameters in our rendering framework that can be tuned to improve on rendering time



Figure 7: Comparison renderings using different BRDFs for the car body in the same lighting. a) Phong. b) Industrial state of the art car paint model (based on the ClearCoat™ model). c) Our measured car paint BRDF. d) Our fitted BRDF model. Note the different distribution of reflection intensity in the Phong example, and the different placement and shape of the highlights in the ClearCoat™ example. Also note that the fitted BRDF matches the tabulated BRDF well and completes the missing data at grazing angles.




Figure 6: Top row: Close up of the sparkle effect, with changing viewpoint. Note the coherence of the sparkles as well as the view dependent shading. Bottom row: Sparkles are less visible and eventually invisible with growing distance.

or quality: Most importantly, the overall rendering time is not dominated by the BRDF model, but by the number of directional lights l used for representing the environmental lighting, the number of samples s for the most specular lobe, and number of samples per pixel p for antialiasing edges and flakes. By varying these parameters we have a fine granular control over rendering speed and quality. Furthermore, if $s = 1$ it may be desirable to always sample the perfect specular direction to avoid the distracting high variance in the estimate at the cost of unrealistic sharp reflections.

Additionally we use the efficient interleaved sampling [13] technique to effectively reduce the number of shadow rays to l/p and the number of reflection rays to s/p . This is achieved by assigning only a subset of the directional lights and the reflection rays to each primary ray of one pixel. As the sampling pattern is only changing at subpixel level no filtering in image space is necessary.

6 Results

For the results presented in the following we used a cluster of 20 dual Opteron 2.5 GHz PCs. If not stated otherwise the images are rendered at video resolution (640×480).

6.1 Comparison to Other BRDF Models

First we compare our acquired car paint BRDF with the Phong BRDF [20], and the industrial car paint model ClearCoat™. Because of its simplicity the Phong model is very popular in computer graphics. However the neglected Fresnel dependency results in physically implausible reflection behavior in particular at grazing angles, as can be seen in Figure 7a.

The more advanced ClearCoat™ BRDF – which is basically an enhanced Phong BRDF – includes a Fresnel term to control the reflection behavior and thus is able to also model the clear coat layer of car paints. Nevertheless, it is not easy and intuitive to manually tune the parameters of this model – the resulting appearance still looks artificially.

However, the combination of a physically based BRDF model together with the measured data from real car paint enables us to create convincing images of virtual cars when rendered using our fitted car paint BRDF (Figure 7d).

6.2 Quality and Performance

We also explore the influence of different parameter sets on quality and rendering performance. Using a moderate discretization of the environment by 16 directional lights, no over-sampling, and one reflection ray allows for realtime performance of over 10 fps. Even at rather high quality settings

($l = 128, s = 16, p = 16$) interactive frames rates of about one frame per second can be maintained and the rendering result is visually pleasing (see Figure 8b): Shadows are smooth, aliasing artifacts are suppressed, and the reflections are slightly blurred. Our framework can be well used to create high quality offline renderings (Figure 8c) at high resolutions within minutes.

7 Conclusion and Future Work

In this paper we have presented a framework for efficient acquisition and rendering of car paint. Our image based acquisition system is easy to build and can measure BRDF of real-world car paint samples quickly and accurately. The acquired BRDF data can be directly used for rendering or it can be used to fit the parameters of a multi-lobe Cook-Torrance BRDF model.

We found that the rendered images – using the measured data in either way – look realistic and outperform renderings using other BRDF models (with often manually tuned parameters) in terms of visual quality (see also Figure 10).

As the BRDF cannot represent the microstructure of metallic car paints that is responsible for the sparkling effect we proposed a fast and frame-to-frame coherent method to simulate these effects.

For future work we like to extend the car paint BRDF database and use it to further analyze car paints. Additionally, a multi-level model for the simulation of the sparkling would be interesting to further reduce the aliasing.

Acknowledgements

We like to thank Christian Fuchs and Thomas Hirth for painting the spheres as well as Martin Fuchs and Rafał Mantiuk for helpful discussions.

References

- [1] S. Agarwal, R. Ramamoorthi, S. Belongie, and H. Wann Jensen. Structured Importance Sampling of Environment Maps. *ACM Trans. Graph.*, 22(3):605–612, 2003.
- [2] J.-Y. Bouguet. Camera Calibration Toolbox for Matlab. www.vision.caltech.edu/bouguetj/calib_doc/.
- [3] B. Cabral, M. Olano, and P. Nemeč. Reflection Space Image Based Rendering. In *SIGGRAPH '99*, pages 165–170, 1999.
- [4] R. L. Cook and K. E. Torrance. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.*, 1(1):7–24, 1982.
- [5] P. E. Debevec. Rendering Synthetic Objects Into Real Scenes: Bridging Traditional and Image-Based Graphics With Global Illumination and High Dynamic Range Photography. In *SIGGRAPH '98*, pages 189–198, 1998.
- [6] K. Devlin, A. Chalmers, A. Wilkie, and W. Purgathofer. Tone Reproduction and Physically Based Spectral Rendering. In *Eurographics State of the Art Reports*, pages 101–123, 2002.
- [7] S. Ershov, K. Kolchin, and K. Myszkowski. Rendering Pearlescent Appearance Based on Paint-Composition Modelling. *Computer Graphics Forum*, 20(3):221–238, 2001.
- [8] S. Ershov, R. Đuriković, K. Kolchin, and K. Myszkowski. Reverse engineering approach to appearance-based design of metallic and pearlescent paints. *The Visual Computer*, 20(8-9):586–600, 2004.
- [9] S. V. Ershov, A. B. Khodulev, and K. V. Kolchin. Simulation of Sparkles in Metallic Paints. In *Proc. Graphicon*, pages 121–128, 1999.
- [10] M. Goesele, H. P. A. Lensch, W. Heidrich, and H.-P. Seidel. Building a Photo Studio for Measurement Purposes. In *Proceedings of VMV'00*, pages 231–238, 2000.
- [11] J. T. Kajiya. The Rendering Equation. In *SIGGRAPH '96*, pages 143–150, 1986.
- [12] C. Kelemen and L. Szirmay-Kalos. A Microfacet Based Coupled Specular-Matte BRDF Model with Importance Sampling. In *Eurographics Short Presentations*, pages 25–34, 2001.
- [13] A. Keller and W. Heidrich. Interleaved Sampling. *EGWR '01*, pages 269–276, 2001.
- [14] T. Kollig and A. Keller. Efficient Illumination by High Dynamic Range Images. In *EGWR '03*, pages 45–50, 2003.
- [15] E. Lafortune. *Mathematical Models and Monte Carlo Algorithms for Physically Based Rendering*. PhD thesis, Katholieke Universiteit Leuven, Belgium, 1996.
- [16] S. R. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, 1998.
- [17] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A Data-driven Reflectance Model. *ACM Trans. Graph.*, 22(3):759–769, 2003.
- [18] M. D. McCool, J. Ang, and A. Ahmad. Homomorphic Factorization of BRDFs for High-performance Rendering. In *SIGGRAPH '01*, pages 171–178, 2001.
- [19] A. Ngan, F. Durand, and W. Matusik. Experimental Validation of Analytical BRDF Models. In *ACM SIGGRAPH: Sketches and Applications*, 2004.
- [20] B. T. Phong. Illumination for Computer Generated Pictures. *Commun. ACM*, 18(6):311–317, 1975.
- [21] W. H. Press, W. T. Vetterling, S. A. Teukolsky, and B. P. Flannery. *Numerical Recipes in C++: the art of scientific computing*. Cambridge University Press, 2002.
- [22] M. A. Robertson, S. Borman, and R. L. Stevenson. Estimation-Theoretic Approach to Dynamic Range Enhancement using Multiple Exposures. *J. Electronic Imaging*, 12(2):219–285, 2003.
- [23] S. Rusinkiewicz. A New Change of Variables for Efficient BRDF Representation. In *EGWR '98*, pages 11–22, 1998.
- [24] R. Đuriković and W. L. Martens. Simulation of Sparkling and Depth Effect in Paints. In *Proc. SCCG*, pages 207–213, 2003.
- [25] E. Veach. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1997.
- [26] I. Wald. *Realtime Ray Tracing and Interactive Global Illumination*. PhD thesis, Saarland University, 2004.
- [27] I. Wald, P. Slusallek, C. Benthin, and M. Wagner. Interactive Rendering with Coherent Ray Tracing. *Computer Graphics Forum*, 20(3):153–164, 2001.



Figure 8: Comparison of renderings with different quality settings generated by 20 dual Opteron 2.5 GHz PCs. Left: Low quality settings (light sources $l = 16$, specular lobe samples $s = 1$, samples per pixel $p = 1$) to gain realistic frame rates of 12.1 fps. Note the visible shadow boundaries, the aliasing artifacts as well as the unrealistic sharp reflections. Center: With medium quality settings ($l = 128, s = 16, p = 16$) most of the aliasing artifacts already vanished. The performance of 1.3 fps still allows for interacting with the model. Right: High quality settings ($l = 1024, s = 256, p = 64$) are used for offline rendering (97 seconds) at a high resolution of 1280×960 .




Figure 9: Our sphere-samples painted with different car paints. Top row: Metallic car paints. From left to right: Ford F8, Polaris Silber, Opel Titan, and BMW 339. Bottom row: Standard car paints. From left to right: 2K Acrylack, White, Blue Matte, and Blue.



Figure 10: Rendering results with different measured and fitted car paints. From left to right, top to bottom: Ford F8, Polaris Silber, Opel Titan, BMW 339, 2K Acrylack, White, Blue Matte, and Blue.