

Measurement and editing of metallic car paint BRDF

Martin Rump *

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

Recent measurement devices for the Bidirectional Texture Function (BTF) offer high parallel capture of the spatial varying reflectance data for material samples. We present an approach to use measurement data from homogenous materials or materials with stochastical mesostructure to obtain data samples for the fitting of BRDF models like Cook-Torrance. We applied this approach to measured samples of metallic car paint and present a novel BRDF model that enables simple editing of the car paint's color by dividing the BTF into a model-based part, a direction-dependent color table and the spatial varying glitter effects from the effect flakes in modern car paints.

1 Introduction

Nowadays image based measurements of the Bidirectional Reflectance Distribution Function (BRDF) and the Bidirectional Texture Function (BTF) are commonplace. But still the measurement devices are different. Most image based BRDF measurement devices use curved material samples to obtain as many data samples per image as possible in contrast to BTF measurement devices that need flat samples. A common assumption when measuring BTFs is orthographic

projection and directional lighting resulting in the same incident and outgoing angles for all Apparent BRDFs (ABRDF) of the BTF. But since the cameras in such devices as well as the light sources have finite distance compared to the size of the material samples, these angles vary across the area of the sample. This enables the extraction of BRDF samples from the BTF measurements and fitting of analytical models to this data samples.

One interesting class of materials with stochastic spatial varying reflectance are modern metallic car paints. They show both significant angular and spatial variation of reflectance and in some cases angular dependent color shifts. We present a modified version of the Cook-Torrance BRDF model that is able to reproduce all relevant effects of modern car paint, which are:

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

Furthermore the model allows for simple changes of the car paints color which enables a designer to visualize a paint in different colors using only one measurement.

*e-mail: rump@cs.uni-bonn.de

2 Previous Work

BRDF measurements In computer graphics literature there is a rich branch about measuring the BRDF for photo-realistic rendering.

The use of measured BRDFs in graphics was introduced by Ward [13] 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. [10] qualitatively compared several popular BRDF models. They used a database of about 100 isotropic BRDFs captured by Matusik et al. [8] which also contains several car paint samples and concluded that physical based models like the Cook-Torrance model [1] are better suited for these kind of materials than empirical models like the Ward [13] or Lafourche [5] models. Furthermore, they exemplified that more than one lobe is needed for these kind of layered materials.

BTF measurements The Bidirectional Texture Function was introduced by Dana et al. in [2]. They presented their measurement setup consisting of a video camera, a halogen lamp and a robot arm to orient the material sample. This device captured 205 images per material sample. They published their BTF database (CURET) consisting of 61 different measured samples from many classes of natural materials. Since then a lot of work was done in the area of BTF measurement. Most devices use a gonioreflectometer setup with CCD-based cameras. For an overview we refer to the State-of-the-Art Report from Müller et al. [9].

Car paint measurements The findings of Ngan et al. [10] had some influence on the work of Guenther et al. [4] who built an BRDF acquisition setup similar to Marschner et al. [7] in order to obtain sampled car paint BRDFs. They fitted a Cook-Torrance [1] 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. [3]. Their work is probably the most related to the approach presented in this paper. The main difference is that we rely on BTF instead of BRDF measurements and do not use any simulation.

Another work on measurement has been done by Sung et al. [12]. They determine individual flake orientations in the car paint by using a confocal laser scanning microscope and gain a distribution



Figure 2: The BTF measurement device consists of 151 digital consumer cameras mounted on a hemispherical gantry.

function from this measurements. Additionally 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 introduce the BTF measurement procedure we used for appearance capture and explain necessary calibration steps. The section 4 introduces our model for car paint and describes in detail how to efficiently extract BRDF samples from BTF data and how to fit a BRDF model to these datapoints. Section 5 describes our car paint editing tool and section 6 the rendering method. We close the paper with results and conclusions.

3 Measurement

We propose to capture the appearance of spatially varying car paints using BTF measurements. Generally any of the standard devices reviewed in [9] 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 2). 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 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.

3.1 Geometrical Calibration

For the analytical model fitting procedure described in Section 4.1 we need an as accurate as possible calibration of the intrinsic and extrinsic parameters of each camera. Since the light sources are a fix part of the cameras, no special calibration procedure is needed for them. of the cameras, we only have to determine position and orientation of the cameras.

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 a rectangular grid of 121 LEDs as features. 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. We use Zhang's camera calibration algorithm [14] to solve for the distinct extrinsic and the common intrinsic parameters and optimize for the individual intrinsic parameters afterwards

3.2 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 (**Flash Quantity**) 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. 136dB, corresponding to an RGB representation with 23bit per channel.

3.3 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 four reflectance standard targets from 2.3% to

96.9% reflectivity that are positioned near the paint sample.

4 Data Preparation

4.1 Modeling

In this section we explain our model for metallic car paints. The basic idea is to split the reflectance in three parts:

- A homogeneous intensity part modelled by a lobe based BRDF model
- An angular dependent color table to model the color flop effects
- A spatial varying part using standard BTF techniques

As suggested by Ngan et al. [10] we use the Cook-Torrance BRDF model [1] with multiple lobes to represent the homogeneous BRDF part of the car paint. The formula is

$$\rho_X(\omega_i, \omega_o) = \frac{k_d}{\pi} + \sum_{i=1}^K \frac{k_{si}}{\pi} \frac{F_{F0,i} D_{m_i} G}{\cos \theta_i \cos \theta_o}. \quad (1)$$

Here ω_i, ω_o are the incoming and outgoing directions. k_d is the diffuse intensity, k_{si} , m_i and $F_{F0,i}$ are the per-lobe specular coefficient, the distribution exponent and the Fresnel parameter respectively. K is the number of lobes, D is Beckmann's microfacet distribution. Both the microfacet distribution D and the geometric attenuation term G are from [1]. For the Fresnel term the approximation of Schlick is used. X denotes the whole set of parameters.

To model the color shift of the car paint 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 measurements the intensity would be

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

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

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

Since our measurements are in linearized RGB we have $N = 3$ and end up with the following formula:

$$I = \frac{R + G + B}{3} \quad (4)$$

This enables dividing color and intensity of the homogeneous BRDF part:

$$\begin{pmatrix} \hat{R} \\ \hat{G} \\ \hat{B} \end{pmatrix} = \frac{1}{I} \cdot \begin{pmatrix} R \\ G \\ B \end{pmatrix} \quad (5)$$

The BRDF model 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). Since this table contains the color information from the paint it is possible to recolor the material by simple substitution of the table. Compressing the table using for example piecewise polynomial basis functions like B-splines should work well and is left for future work.

To obtain the color table χ we simply average the rectified images of the BTF datasets and store the resulting color together with the corresponding incident and outgoing direction in the table. During evaluation linear interpolation of the data points is used.

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

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

The last lobe models 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.

4.2 BRDF Sample extraction

In this chapter we describe our algorithm to extract BRDF data samples from the BTF measurement data.

We rely on the fact that the cameras of the measurement device have about 64cm distance from a 10x10cm planar material sample. This induces an angular variation of about 12° from one corner of the sample to the opposite corner.

The basic idea of the data sample extraction is to apply an adaptive box filter to the rectified BTF images to filter out the spatial variation introduced by the flakes. This is simply done by choosing a box size and averaging over image patches of this size. The exact light- and view directions are calculated for the center of the patch based on the calibration data from the calibration step described in Section 3.1. The patch average and the two directions are one sample for the BRDF.

We noticed during our experiments that fitting of the whole BRDF model to these samples at once is numerically unstable (what was also reported by Günther et al. [4]). The optimizer fails to fit the gloss lobe because there are too few samples of this lobe in the whole sampleset. To solve this problem, we extract two different sets of samples from the BTF data. The first set Ψ_d represents the diffuse and glitter part of the BRDF while the second set Ψ_s contains data for the gloss lobe.

To extract the diffuse/glitter sampleset Ψ_d 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: $p := w_{low} + \sqrt{\theta_H/\theta_{H,max}}(w_{high} - w_{low})$. (θ_H denotes the angle between halfway vector and normal vector, $\theta_{H,max}$ is the greatest accepted value for θ_H of a sample.)
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.

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 .

The second sampleset Ψ_d is chosen only from images which contain the mirror direction and are very accurately calibrated. The quality of the calibration is measured 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 filter for this second sampleset must be chosen carefully since it must be as small as possible in order to preserve as much information about the shape of the gloss lobe as possible but 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) lead to the best results.

4.3 The fitting procedure

Having extracted enough samples an optimization process is started to find the parameter set X that leads to the smallest error when comparing the data samples and the BRDF model evaluated with the same angles. Thus the optimization process minimizes an error function well suited for this problem by iteratively changing the BRDF parameters X .

We use an error function that is similar to the one from Günther et al. [4]:

$$E(X) = \sum_{(I, \omega_i, \omega_o) \in \Psi} \left| \log \frac{I + \epsilon}{M(X, \omega_i, \omega_o) + \epsilon} \right| + P(X) \quad (7)$$

Here Ψ is one of the sets of data samples from the previous step, M is the Cook-Torrance BRDF model with parameters X . A small positive constant ϵ prevents division by 0, since for some parameters X the BRDF model may be 0. P is a plausibility function for the parameters X and depends on the BRDF model:

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

With a big constant K .

We implemented a quasi-newton-minimizer using the SR-1 update method for the hessian matrix. To overcome problems with local minima of the error function we start the fitting process from several start values and take the best result.

The following pseudo-code describes the whole BRDF fitting algorithm. Parameter sets X are divided into $X_s = (k_l, m_l, F_{0l})$ for the last lobe modeling the surface gloss and X_d for the rest of the BRDF. Parameter sets with subscript 0 are start values.

```
for j=1 to J do
     $X_{j,d} = \text{Optimize}_d(\Psi_d, M, E, X_{0,j,d}, X_{0,s})$ 
```

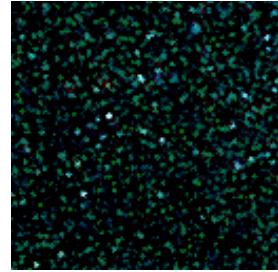


Figure 3: One image from the difference \overline{BTF}

```
next
 $X_d = X_{j,d}$  mit  $E(X_{j,d}, X_{0,s})$  minimal
for j=1 to K do
     $X_{j,s} = \text{Optimize}_s(\Psi_s, M, E, X_d, X_{0,j,s})$ 
next
 $X_s = X_{j,s}$  mit  $E(X_d, X_{j,s})$  minimal
for j=1 to K do
     $X_d = \text{Optimize}_d(\Psi_d, M, E, X_d, X_s)$ 
     $X_s = \text{Optimize}_s(\Psi_s, M, E, X_d, X_s)$ 
next
```

In the first part J resp. K start values X_0 are tried to gain parameters for X_d, X_s . The best result from X_d is already taken for the fitting of X_s . In the second part the parameters are corrected by turns to compensate for the error induced by the separate fit process.

4.4 Separation of BRDF and BTF part

To separate the homogenous BRDF part from the spatial varying BTF part we simply evaluate the BRDF for every pixel of every rectified BTF image and make a simple subtraction in RGB space. The light and view directions for the pixels can be calculated from the camera calibration data from Section 3. The resulting difference images form a new BTF which we simply call $\overline{BTF}(x, \omega_i, \omega_o)$:

$$\overline{BTF}(x, \omega_i, \omega_o) = BTF(x, \omega_i, \omega_o) - \rho(\omega_i, \omega_o) \quad (9)$$

Figure 3 shows one images off such a difference BTF. As it can be seen the base color of the paint is no longer part of the BTF. It does only contain colorful effects from the coated flakes. This enables the replacement of the paints color by modifying the color table χ only.

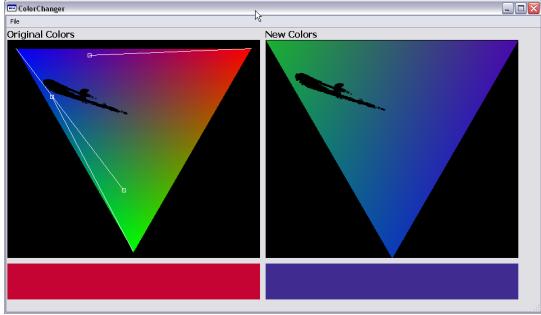


Figure 4: The Color Changer tool. The user can drag the primary valences to new positions creating an affine mapping of the color plane

5 Color Editing

Since the color information of the car paint resides in our color table χ we are able to offer very simple editing of the paint. We created a tool that enables the user to define an affine 2D transformation that maps the original colors in the table to new ones. This new table can then directly be used for rendering. Figure 4 shows a screenshot of the tool in use.

The transformation T is defined as solution of the following equation:

$$\begin{aligned} TS &= D \\ T &= \begin{bmatrix} m_{11} & m_{12} & t_1 \\ m_{21} & m_{22} & t_2 \\ 0 & 0 & 1 \end{bmatrix} \\ S &= \begin{bmatrix} s_{1x} & s_{2x} & s_{3x} \\ s_{1y} & s_{2y} & s_{3y} \\ 1 & 1 & 1 \end{bmatrix} \\ D &= \begin{bmatrix} d_{1x} & d_{2x} & d_{3x} \\ d_{1y} & d_{2y} & d_{3y} \\ 1 & 1 & 1 \end{bmatrix} \quad (10) \end{aligned}$$

The s_i are the source colors and the d_i the destination colors. T maps the d_i onto the s_i .

6 Synthesis and Rendering

As explained in Section 4.1 we need to put together the BTF and the BRDF part of our car paint at rendering time. This is our final spatial varying BRDF:

$$\begin{aligned} \rho(x, \omega_i, \omega_o) &= \overline{BTF}(x, \omega_i, \omega_o) + \\ \chi(\omega_i, \omega_o) &\left(\frac{k_d}{\pi} + \sum_{i=1}^{K-1} \frac{k_{s_i}}{\pi} \frac{F_{F0,i} D_{m_i} G}{\cos \theta_i \cos \theta_o} \right) \\ &+ \frac{k_{s_K}}{\pi} \frac{F_{F0,K} D_{m_K} G}{\cos \theta_i \cos \theta_o} \end{aligned}$$

This model 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 size of the uncompressed data is too large for effective rendering. 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 [11] 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).

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 [6] 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.

7 Results

We applied the data preparation and fitting process described in Section 4 to 3 different car paints which are depicted in Figure 6. 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 5. It is particular since it contains direct sunlight. The synthesis and rendering stage described in Section 6 was implemented as a MentalRay-shader. Furthermore we recolored the car paints and recombined the BRDF and BTF



Figure 5: The main entrance environment map

parts to create new paints. Results can be seen in Figure 7.

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. Further dividing the BRDF model into an intensity and a color part enables for simple editing of the car paints color.

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. Furthermore, we want to integrate our model within a PRT rendering framework in order to achieve real-time rendering with global illumination effects.

References

- [1] Robert L. Cook and Kenneth E. Torrance. A reflectance model for computer graphics. In *SIGGRAPH '81: Proceedings of the 8th annual conference on Computer graphics and interactive techniques*, pages 307–316, New York, NY, USA, 1981. ACM Press.
- [2] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 151–157, 1997.
- [3] Sergey Ershov, Andrei Khodulev, and Konstantin Kolchin. Simulation of sparkles in metallic paints. In *Proceedings of Graphicon '99*, pages 121–128. The Eurographics Association and Blackwell Publishers, 1999.
- [4] Johannes Günther, Tongbo Chen, Michael Goesele, Ingo Wald, and Hans-Peter Seidel. Efficient acquisition and realistic rendering of car paint. In Günther Greiner, Joachim Hornegger, Heinrich Niemann, and Marc Stamminger, editors, *Vision, Modeling, and Visualization 2005 (VMV'05)*, pages 487–494, Erlangen, Germany, November 2005. Aka.
- [5] Eric P. F. Lafourte, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear approximation of reflectance functions. In *Proceedings of SIGGRAPH*, pages 117–126. ACM Press/Addison-Wesley Publishing Co., 1997.
- [6] Jason Lawrence, Szymon Rusinkiewicz, and Ravi Ramamoorthi. Adaptive numerical cumulative distribution functions for efficient importance sampling. In *Eurographics Symposium on Rendering*, June 2005.
- [7] Stephen R. Marschner, Stephen H. Westin, Eric P. F. Lafourte, Kenneth E. Torrance, and Donald P. Greenberg. Image-based BRDF measurement including human skin. In *Proceedings of 10th Eurographics Workshop on Rendering*, pages 139–152, 1999.
- [8] Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. Efficient isotropic BRDF measurement. In *Proceedings of the EGRW '03*, pages 241–247, 2003.
- [9] G. Müller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein. Acquisition, synthesis and rendering of bidirectional texture functions. *Computer Graphics Forum*, 24(1):83–109, March 2005.
- [10] Addy Ngan, Frédéric Durand, and Wojciech Matusik. Experimental analysis of brdf models. In *Proceedings of the Eurographics Symposium on Rendering*, pages 117–226. Eurographics Association, 2005.
- [11] M. Sattler, R. Sarlette, and R. Klein. Efficient and realistic visualization of cloth. *Proceedings of the Eurographics Symposium on Rendering 2003*, 2003.
- [12] L. P. Sung, M. E. Nadal, M. E. McKnight, E. Marx, R. Dutruc-Rosset, and B. Laurenti. Effect of aluminum flake orientation on coating appearance. *Proceedings of the 79th Federation of Societies for Coatings Technology (FSCT ICE) Meeting*, pages 1–15, nov 2001.
- [13] Gregory J. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of SIGGRAPH*, pages 265–272. ACM Press, 1992.
- [14] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11):1330–1334, 2000.

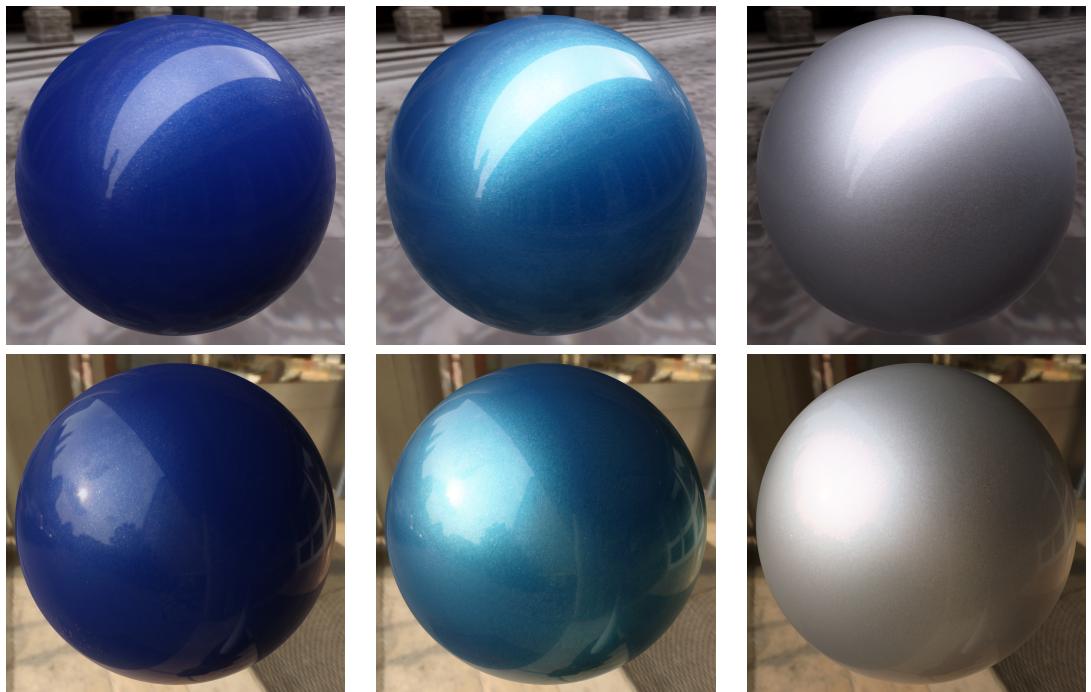


Figure 6: The original car paints rendered in two different environments. Top row: Uffizi environment. Bottom row: Main entrance environment



Figure 7: Results of car paint editing. Left: Darkblue paint with other colors. Middle: Green-blue paint with other colors and flakes from the silver paint. Right: Silver paint with same colors but flakes from green-blue paint.