

Attribute-preserving gamut mapping of measured BRDFs

Tiancheng Sun^{1,2} Ana Serrano² Diego Gutierrez² Belen Masia²¹ IIS, Tsinghua University ² Universidad de Zaragoza, I3A

Abstract

Reproducing the appearance of real-world materials using current printing technology is problematic. The reduced number of inks available define the printer's limited gamut, creating distortions in the printed appearance that are hard to control. Gamut mapping refers to the process of bringing an out-of-gamut material appearance into the printer's gamut, while minimizing such distortions as much as possible. We present a novel two-step gamut mapping algorithm that allows users to specify which perceptual attribute of the original material they want to preserve (such as brightness, or roughness). In the first step, we work in the low-dimensional intuitive appearance space recently proposed by Serrano et al. [SGM*16], and adjust achromatic reflectance via an objective function that strives to preserve certain attributes. From such intermediate representation, we then perform an image-based optimization including color information, to bring the BRDF into gamut. We show, both objectively and through a user study, how our method yields superior results compared to the state of the art, with the additional advantage that the user can specify which visual attributes need to be preserved. Moreover, we show how this approach can also be used for attribute-preserving material editing.

CCS Concepts

•Computing methodologies → Reflectance modeling; Perception;

1. Introduction

Real-world materials present a wide variety of appearances, commonly described in computer graphics with the bidirectional reflectance distribution function (BRDF). Current printers, on the other hand, have a predefined set of only a few inks, which in turn defines the printer's gamut. As a consequence of this limitation, many materials cannot be exactly reproduced by the printer. Finding the best approximation of the input BRDF that falls within the printer's gamut is the problem known as BRDF gamut mapping.

Gamut mapping is an extremely underconstrained problem without a unique solution, for which several methods have been proposed. The usual approach is to try to find the BRDF which is *most similar* to the target BRDF, while lying within the available gamut. In this paper, we add a novel approach to the state of the art: Our gamut mapping technique allows the user to set a certain perceptual attribute (or attributes) that needs to be preserved as much as possible while mapping the BRDF into the available gamut. For instance, the user may specify explicitly the preservation of the *strength of reflections*, or the *metallic* appearance of the material.

To achieve our goal, we propose a two-step gamut mapping technique: In the first step, we leverage recent works on material acquisition [NJR15] and editing [SGM*16]. In these works, Nielsen et al. first built a five-dimensional PC space which serves as a basis for representing each BRDF; then, Serrano et al. learnt functionals mapping the space of principal components to high

level perceptual attributes defined for achromatic reflectance; these functionals define an intuitive control space for appearance. We use these mappings in PC space to follow the path that brings the luminance channel L (L in Lab space) into gamut in PC space, while preserving the desired attribute (see Figure 1). However, adding the *ab* color coordinates to the remapped *L* leads to a BRDF that will most likely still be out of gamut (Figure 1). We thus complete the gamut mapping process with our second step, which consists of an image-based optimization, inspired by other recent works [PR12, NDM06].

We validate our technique on the BRDFs from the MERL database [MPBM03], using a gamut formed by real measured inks [MAG*09]. We show that, for a majority of BRDFs in the database, performing this first step we propose leads to a better final result than that obtained by previous works. Given that a definitive metric for BRDF similarity does not exist, in addition to computing a BRDF metric, we validate our results by means of a user study. Although it is not a requirement of the method, our proposed method allows certain interactivity, since the user can choose which attribute(s) to preserve. In the following, when performing comparisons to the state-of-the-art automatic method, we always fix the same two attributes, in order to provide a fair comparison to it. However, we also show results preserving other attributes, showing this additional feature of our method.

2. Related work

In this section we focus on perceptual spaces for BRDFs and gamut mapping works; we refer the reader to more general recent surveys on perception and graphics [MMG11] and BRDF representation [GGG*16] for a broader view.

2.1. Perceptual spaces for BRDFs

Predicting the appearance of a given BRDF can be complicated: not only do the shape of the object and the lighting environment affect the perceived appearance of the material [VLD07,FDA03], but directly changing the parameters of analytical models usually results in unintuitive and perceptually non-uniform changes as well. Observing this, Pellacini *et al.* derived a perceptually uniform parameter space for the Ward BRDF model [PFG00]. Ngan *et al.* proposed an image-driven BRDF metric, which allows users to linearly modify the material's appearance using several analytical models. A perceptual space for gloss was later proposed based on real material measurements [WAKB09]. Havran *et al.* pushed forward perceptually-motivated BRDFs by further considering the lighting and viewing directions in a single image [HFM16]. Matusik and colleagues presented the MERL database of measured BRDFs [MPBM03], proposed two techniques for dimensionality reduction, and had a single user classify whether a given BRDF possesses a series of *perceptual traits* or not, which in turn allowed them to modify a BRDF's appearance intuitively. Recently, Serrano *et al.* [SGM*16] extended the MERL dataset and presented a series of functionals connecting such data with the perceptual ratings obtained from large-scale user studies. The derivation of these perceptual spaces can be used for gamut mapping, in particular for the establishment of distances among BRDFs which is required for gamut mapping. In this work, given the underconstrained nature of gamut mapping, we take the approach of preserving certain aspects of appearance, and in particular those defined by one or more perceptual attributes. Therefore, we use the functionals derived by Serrano *et al.* [SGM*16] to map the original BRDF to the nearest material inside the gamut that preserves certain perceptual attributes.

2.2. Gamut mapping for BRDFs

Several approaches have been proposed to reproduce specific material appearances. Weyrich *et al.* [WPMR09] used a grid of tilted small facets to achieve custom reflectances; Matusik *et al.* [MAG*09] used halftoning to print a certain appearance on paper using a regular printer. There are also works focusing on reproducing subsurface scattering [HFM*10,DWP*10]. A model which can cover both isotropic and anisotropic appearances was proposed by Lan *et al.* [LDPT13], which combines the micro-facet model with BRDF printing. Closer to our approach, other works focus on finding the closest material inside the valid gamut of a printer. The metric on half-angle curves of the materials was used to resolve the best components of the inks in the work of Matusik *et al.* [MAG*09], while Lan *et al.* [LDPT13] calculated the L2 norm on the BRDF hemisphere for data fitting. Pereira and Rusinkiewicz improved this procedure by minimizing the difference between rendered images of the original and the final materials [PR12]; one

slice of the BRDF image was calculated analytically for optimization. Similar to this work, we also perform an image-based optimization; however, we do this as a second step in our gamut mapping algorithm, after finding an intermediate BRDF that better preserves the desired perceptual attributes of the original BRDF.

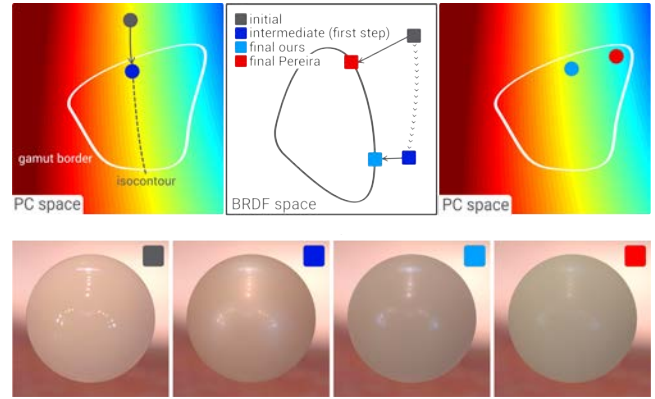


Figure 1: Overview of our two-step gamut mapping method. **Top-left:** A two-dimensional projection of the five-dimensional PC space. The white line represents the border of the gamut, and same color-coded isocontours indicate the same value of a given perceptual attribute (following [SGM*16]). Working on achromatic reflectance, we first push the original (target) BRDF (gray) into gamut (intermediate BRDF, dark blue) in such PC space. **Top-middle:** Back in the original BRDF space, the intermediate BRDF is not guaranteed to be in gamut (the dotted line represents the previous move in PC space); we therefore apply an image-based optimization to bring it into gamut (final BRDF, light blue). The red dot represents the result of applying a single step based on image optimization [PR12]. **Top-right:** For comparison purposes, we move back to PC space to show the final BRDFs with both methods; ours (light blue) lies much closer to the intended attribute value than the single-step method (red). **Bottom:** Real results with the alumina oxide BRDF from the MERL database. From left to right: original (out of gamut); our intermediate BRDF (still out of gamut); our final result; single-step image-based optimization [PR12]. Our result preserves highlights better, while exhibiting less color shift.

3. Attribute-preserving gamut mapping

Our goal is to take an out-of-gamut BRDF and bring it into a representable gamut, defined for instance by the individual color inks of a printer, while preserving a given perceptual attribute, such as its *brightness*. However, mapping perceptual attributes of a material to its underlying BRDF representation is not an obvious task. Recently, Serrano *et al.* [SGM*16] gathered a large number of subjective ratings on the high-level perceptual attributes that best described their extended MERL BRDF dataset. Using these ratings, the authors then built and trained radial basis functions networks, which are used as functionals mapping the perceptual attributes to a five-dimensional log-mapped PCA-based representation of the BRDFs, proposed by Nielsen and colleagues [NJR15]. These functionals, derived for achromatic reflectance only, are good pre-

dictors of the perceived attributes of appearance. We leverage this work to develop our two-step gamut mapping algorithm.

Figure 1 qualitatively presents an overview of our method. First (Figure 1, top left), working in PC space, we follow the isocontour of a given functional to bring into gamut (in PC space) the initial BRDF; these isocontours represent the same value of a given attribute (please refer to [SGM*16] for details). For visualization purposes we show a 2D slice of the original 5D space. This implies fixing values for the other three dimensions to select which slice to plot, and projecting the points corresponding to the BRDFs onto the slice. As a result, points in a particular projection may appear slightly inside the gamut, even though they do lie on the border in 5D space. Note that in this space we only work with L values; color will be handled in a second step. This yields an intermediate BRDF which, although it preserves the desired perceptual attribute, cannot be guaranteed to fall within the gamut defined by the inks in the original BRDF space. In our second step, we bring the intermediate BRDF into gamut using an image-based optimization (Figure 1, top center). Figure 1, top right, shows the final BRDF using our method, compared to a single-step, image-optimization method (such as Pereira’s state-of-the-art algorithm [PR12]) in PC space. It can be seen how our result better preserves the intended attribute in this space, since it is never explicitly taken into account in existing single-step methods. Figure 1, bottom, shows real examples produced with our method, and Pereira’s [PR12]. We use the *alumina oxide* BRDF from the MERL database, and aim to preserve the *metallic-like* and *bright* attributes. Although obvious differences exist in both results with respect to the original BRDF, given the limited ink gamut, our method maintains stronger highlights and exhibits significantly less color shift.

Note that in our second step, we optimize both the L and the chromatic coordinates (a, b) , despite the fact that L had already been modified in the previous step. This is because our gamut is defined by a series of real-world inks, in which L cannot be isolated and optimized independently of chromaticity. In other words, although decoupling achromatic reflectance from color information was convenient in PC space to follow an isocontour that would preserve the value of a certain attribute, they cannot be decoupled when handling real inks. Modifying such ink components independently would not guarantee that the final BRDF lies within the available gamut. Our first step provides, however, a better starting point for the image-space optimization step, leading to better results (see Section 4). In the rest of the section we describe our two steps in detail.

3.1. Step 1: Luminance mapping in PC space

We call our initial BRDF ρ_{ini} . In this first step, our goal is to obtain an intermediate BRDF ρ_{int} , where we bring the L channel into gamut in log-mapped PC space [NJR15] following the path that maintains the same value v_A of a certain perceptual attribute A^\dagger . Since the functionals derived in the work of Serrano et al. are learned with respect to the coefficients for L in PC space, we follow

[†] For the sake of clarity, we explain our method for the simpler case of fixing just one attribute.

here the same procedure. First, we apply a log-relative mapping to ρ_{ini} , which enables a good distribution of the available dynamic range [MAA01, NJR15, SGM*16], and obtain the first five coefficients $\alpha_{\text{ini}} \in \mathbb{R}^5$ of the BRDF in the PCA basis (which provide general hints about material appearance [NJR15]). Then, we apply a function $f: \mathbb{R}^5 \rightarrow \mathbb{R}$ that maps a BRDF in the aforementioned 5D PC space to its attribute value [SGM*16]: $v_A = f(\alpha_{\text{ini}})$. This function f is a radial basis function network (RBFN) based on Gaussian functions such that:

$$f(\alpha) = \sum_{i=1}^{N_c} \theta_i \exp^{-\beta \|\alpha - c_i\|^2} \quad (1)$$

where β controls the smoothness of the Gaussian functions, N_c is the number of neurons in the network, c_i are the centers of such neurons, and θ_i are the weights of each neuron. RBFNs are a particular type of artificial neural networks called static neural networks, where the outputs are linear combinations of radial basis functions, and the neurons correspond to cluster centers.

For our gamut mapping, we formulate the objective function to maintain the initial attribute value $v_A = f(\alpha_{\text{ini}})$, so that the optimization moves along the corresponding isocontour of f as much as possible. Our objective function is therefore:

$$g(\alpha) = f(\alpha) - f(\alpha_{\text{ini}}) \quad (2)$$

Further, we need to ensure that the resulting BRDF is inside the gamut in PC space, which we formulate as a hard constraint. We define the gamut as the set of possible convex combinations of the inks, expressed in our formulation as the convex hull $\text{Conv}(\alpha_{\text{inks}})$ limited by the ink coefficients in the 5D PC space. The resulting optimization problem becomes:

$$\min_{\alpha} \|g(\alpha)\|^2 \quad \text{s.t.} \quad \alpha \in \text{Conv}(\alpha_{\text{inks}}) \quad (3)$$

In order to solve this optimization we use sequential quadratic programming (SQP) [WN99] as implemented in the *fmincon* function in *MATLAB*. In this way, we obtain the new PC coefficients α defining our intermediate BRDF ρ_{int} , which lies inside the inks gamut in PC space, while keeping the value v_A of the desired attribute A from the initial BRDF ρ_{ini} .

If instead of one we want to preserve multiple attributes, we employ a linear combination of them in the objective function. In principle, we weight all attributes equally in this linear combination, but alternative weighting according to the user’s intent is also possible.

3.2. Step 2: Image-space optimization

After the first step we have an intermediate BRDF ρ_{int} which is not necessarily inside the gamut defined by the inks, since we have optimized for achromatic reflectance L only in log-mapped PC space. Let us consider a gamut defined by a set of N inks; any reproducible BRDF lies inside the convex hull formed by the inks’ BRDFs $\mathbf{P}_{\text{inks}} = [\rho_1 \ \rho_2 \ \cdots \ \rho_N]$ in BRDF space. Our goal is to find a BRDF ρ_{fin} that lies within this convex hull, and is thus a convex combination of \mathbf{P}_{inks} such that:

$$\rho_{\text{fin}} = [\rho_1 \ \rho_2 \ \cdots \ \rho_N] \cdot \mathbf{w}, \quad (4)$$

where $\mathbf{w} \in \mathbb{R}^N$ is the vector formed by the coefficients of the convex combination.

When trying to minimize the distance between the intermediate BRDF ρ_{int} and the reproducible one ρ_{fin} , we can in principle work in BRDF space, or in image space. Working in BRDF space with measured BRDFs is very costly, due to the large size of the data; furthermore, similarity of the raw BRDF data does not imply visual similarity [FFGZ12]. Therefore, we will instead minimize this distance in image space, as has been done in the past [PR12].

Given the incident light, geometry, and BRDF, the per-pixel image formation model is described by the rendering equation:

$$L_{\text{out}} = L_{\text{emit}} + \int_{\Omega} \rho(\omega_i, \omega_o) L_{\text{in}} \cdot (\omega_i \cdot \mathbf{n}) d\omega_i. \quad (5)$$

where L_{out} is the outgoing radiance in direction ω_o , L_{in} is the incident light in direction ω_i , \mathbf{n} is the surface normal, and Ω is the hemisphere defining the integral domain; we consider additional light emitted by each surface point L_{emit} to be zero. Given inter-reflections and indirect lighting, this equation defines a non-linear process. However, if we consider a single convex object (a purely opaque sphere, commonly used to visualize BRDFs) light interacts only once on its surface before reaching the camera. Thus, the rendered image and the BRDF are linearly related as:

$$\mathbf{I} = \mathbf{R} \cdot \rho, \quad (6)$$

where \mathbf{R} is a matrix defining the linear mapping. Using Equation 6, we can change Equation 4 into:

$$\mathbf{I} = [\mathbf{I}_1 \quad \mathbf{I}_2 \quad \cdots \quad \mathbf{I}_N] \cdot \mathbf{w}, \quad (7)$$

where \mathbf{I}_i are the rendered images for each ink. Note that this side-steps the need to explicitly compute \mathbf{R} ; moreover, these rendered images allow to establish visual similarity better than raw BRDF data. We can now use these rendered images to obtain the optimal coefficients \mathbf{w}_{opt} :

$$\mathbf{w}_{\text{opt}} = \underset{\mathbf{w}}{\text{argmin}} d(\mathbf{I}_{\text{int}}, [\mathbf{I}_1 \quad \mathbf{I}_2 \quad \cdots \quad \mathbf{I}_N] \cdot \mathbf{w})$$

s.t. $\sum \mathbf{w} = 1, 0 \leq \mathbf{w} \leq 1.$

where \mathbf{I}_{int} is the image obtained with the BRDF computed in the first step. This distance d is thus computed in image space, and could be done in RGB or in Lab color space. We choose d to be the L_2 norm under Lab color space, since it better preserves the color of the original BRDF; we show an example of this in Figure 2.

4. Results

When computing our results, we use the BRDF gamut from Matusik et al. [MAG*09]. This gamut contains 57 BRDFs, which are measured from real world inks. Since the gamut is designed to reproduce a wide range of material appearances, most of the inks are specular and metallic, which are not found in standard printers (the inks are shown in Figure 3). Our images of the inks and initial BRDF used in the optimization are rendered with the *St. Peter's Basilica* environment map, while for the results used in our user study (Section 4.2) we use the *Eucalyptus Grove* map[‡], given that

[‡] Both environment maps are from the Light Probe Image Gallery [Deb98].




Figure 2: Gamut mapped results for the pinkjasper MERL BRDF (middle) optimizing in Lab (left) and RGB (right). Image-based optimization in Lab space better preserves chromaticity. The set of inks that define the gamut can be seen in Figure 3.




Figure 3: All the BRDFs from the gamut provided by [MAG*09], which are measured from real inks.

these illuminations facilitate material perception [FDA03], and in order to employ different illuminations for validation and optimization. In all results shown in this section, except when noted otherwise, we are performing gamut mapping preserving the *metallic-like* and the *bright* attributes in the first step. All the images have been equally tonemapped adjusting exposure and gamma.

Figure 4 shows the influence of our first step (luminance optimization in PC space) in the final results, as opposed to using only the image-based optimization of the second step. Our final result (two steps) is much closer in appearance to the target, out-of-gamut BRDF (*twolayergold*) than the result of a single-step image-based optimization (i.e., without the first attribute-preserving step). The effect of the first step (although in this case compared to the image-based optimization of Pereira [PR12]) can also be seen in Figure 1.

4.1. BRDF gamut mapping results

We compare our results with those from Pereira [PR12] on 94 homogeneous materials from the MERL database [MPBM03]. Some representative results rendered with different environment maps are shown in Figure 5 (please refer to the supplemental material for the complete set). To provide an objective comparison to the state of the art [PR12], we use the cube root cosine weighted RMS metric, which has been reported to perform better than RMS for




Figure 4: Our two-step, attribute-preserving gamut mapping, compared to a single-step optimization. Note how our method helps preserving the specular highlights and overall appearance better.




Figure 5: Comparison of our results and the state-of-the-art method by Pereira [PR12] using different illuminations (Eucalyptus and Grace). In general, our method minimizes color shifts (chrome and goldmetallicpaint3), while better preserving highlights and specular behavior (grayplastic and whitemarble). For very diffuse materials (greenfabric) neither method succeeds due to the specular nature of the inks used (see Figure 3).




Figure 6: Difference between the error of the state-of-the-art image-based optimization [PR12] and that of our method; the error is computed as the cube root cosine weighted RMS [FFGZ12]. Blue indicates better results with our method.

BRDFs [FFGZ12]. This metric is described as:

$$E = \sqrt{\frac{\sum_n ((\rho_{\text{fin}} \cos \theta_i - \rho_{\text{ini}} \cos \theta_i)^2)^{1/3}}{n}} \quad (8)$$

where θ_i is the cosine of the angle between the incident light and the normal. Results of this metric are shown for all MERL BRDFs in Figure 6. We plot the difference between the error of both methods, sorted by increasing values, where blue indicates better results with our method (our error is lower) and red the opposite (our error is higher). Although we do not outperform the single-step method of Pereira [PR12] for all BRDFs in the database, we do in a majority of cases. The user study in the next subsection confirms this.

4.2. User study

Additionally, we have carried out a perceptual study to evaluate the results of our gamut mapping algorithm, with the same gamut used for the previous objective evaluation. We used a subset of 50 out-of-gamut BRDFs from the MERL dataset, discarding in-gamut BRDFs and BRDFs lying very far away from the gamut (see Section 5). Similar to previous works [PFG00, FDA03, PR12] we use a sphere to depict the materials. The design of our user study is based on the one reported by Pereira and Rusinkiewicz, in order to provide a fair comparison. We render them under the *Eucalyptus* environment map. In our study the user is presented with a reference image (center), and two different results (Pereira's and ours), one at each side (see Figure 7). The order in which the BRDFs appeared, as well as the position of each result relative to the ground truth, was randomized. Subjects were asked to select which of the two alternatives shown was more visually similar to

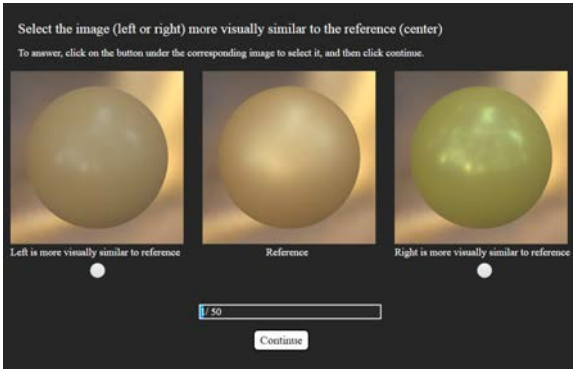


Figure 7: Screenshot of our user study. The reference is presented in the middle, with the two options at the sides in random order.




Figure 8: Vote counts indicating preference for the BRDFs mapped with our method (blue bars) and Pereira’s method (red bars) for the BRDFs with statistically significant differences. In them, our result was preferred 17 out of 22 times with high agreement between users.

the reference image. They were instructed that by *visually similar* we meant which of the two better represented the material appearance of the ground truth sphere.

We recruited fifteen subjects (nine male, six female). All subjects were presented with all 50 BRDFs, and the time to completion of the experiment was approximately 10 minutes. There was no time limit for making each choice, but once subjects moved forward to the next example, they were not allowed to go back. Twenty-two of the tested BRDFs showed significant differences in the results (χ^2 test, $p < 0.05$); among these, 77.6% of the time our result was chosen over the state-of-the-art method (see Figure 8). Agreement between users was high, with 81.3% users on average agreeing with the majority on a given choice. Overall, including the non-statistically significant BRDFs, our results were preferred in almost 62% of the results.

4.3. Additional results

In this section we present additional gamut mapping results preserving different combinations of attributes during the optimization along isocontours in our first step. In Figure 9 we show the outcome of using single attributes as opposed to a combination of several attributes. For the cases where the BRDF is far from the




Figure 9: From left to right: Original BRDF and corresponding gamut mapping results computed by preserving, during the first step of our method, only the metallic-like attribute, only the bright attribute, and both attributes. Optimizing over the metallic-like isocontour yields more accurate reflections, while if we optimize over the bright isocontour the diffuse behavior is better preserved. A combination of the two attributes reaches a compromise, aiming to preserve both behaviors.

gamut, it is challenging to obtain a convincing result that matches the appearance of the original BRDF. In such cases, compromises in appearance need to be made in order to obtain a BRDF that can be represented in the gamut. When optimizing to preserve only the *metallic-like* attribute, the resulting BRDF preserves the specular behavior better, while when optimizing for the *bright* attribute, the result matches the diffuse component better. When optimizing with the two attributes at the same time, the optimization reaches a compromise between both. In every case our algorithm presents a predictable behavior, and can be adapted to the user’s needs. In Figure 10 we show additional results preserving in this case the attributes *rough* and *strength of reflections*. Note that this combination of attributes performs particularly well at preserving the look of the reflections, even for BRDFs which are very far away from the gamut.

4.4. Material editing

Our strategy to preserve attributes can also be applied to intuitive material editing, extending the capabilities of the recent work by Serrano et al. [SGM*16]. In particular, we allow the user to adjust the value of one attribute, while fixing a certain value for another one. Let $f^{att_1}(\alpha)$ and $f^{att_2}(\alpha)$ be the functionals in PC space related to the attribute to be changed, and the attribute to be fixed to a given value, respectively. Similar to the original work, we optimize $f^{att_1}(\alpha)$ looking for the α values that yield the desired attribute value y_{obj} . However, we now impose an additional hard constraint over the second attribute, effectively fixing its value y_{fix} :

$$\min_{\alpha} \|f^{att_1}(\alpha) - y_{obj}\|^2 \quad \text{s.t.} \quad f^{att_2}(\alpha) = y_{fix} \quad (9)$$

This provides more accurate control over editing, and can be used to successfully perform edits over attributes originally coupled in Serrano’s work. Figure 11 shows an example where we turn the original *yellowphenolic* BRDF more *metallic-like*, but making the




Figure 10: Results computed preserving both the rough and strength of reflections attributes in the first step of the method, and comparison to the state of the art [PR12]. This particular combination of attributes aims to better preserve the appearance of the reflections.



Figure 11: Extended capabilities for intuitive material editing, turning the original material more metallic. From left to right: the original BRDF; result from Serrano et al. [SGM^{*} 16] (since the metallic-like and bright attributes are coupled in their implementation, the user has no control over the final brightness of the material); our result fixing a low value for bright; and our result fixing a high value for bright.

result more or less bright by fixing different values of the *bright* attribute.

5. Discussion and conclusion

In this paper, we have proposed a new two-step method for BRDF gamut mapping. In the first step we work in PC space, and use some previously proposed functionals that map this space to higher-level intuitive attributes to preserve the appearance of any of such attributes. The output of this first step, which only optimizes achromatic reflectance, is then used as input to an image-space optimization which brings the final BRDF into the ink gamut by expressing it as a convex combination of the available inks. We perform both an objective and subjective validation comparing against the state of the art. Additionally, we show how a slight modification of our framework can provide extended functionalities for intuitive material editing.

We have found out empirically that using the *metallic-like* and *bright* attributes (equally weighted) leads to good results for a large part of the MERL database (please refer also to the supplemental).




Figure 12: Limitations of current methods. If the BRDF lies very far away from the gamut (such as specularmaroonphenolic shown here) both our method and the single-step state of the art are unable to find a satisfying match in appearance. Here, our method does a reasonable job at preserving the specular behavior, but fails to accurately reproduce the diffuse color.

This finding could be used to design an *automatic* gamut mapping method, since *metallic-like* tends to preserve specularities, while *bright* tends to preserve the diffuse color. Apart from these two attributes, users could then tune the weight of other individual attributes, to obtain different results targeted to specific purposes.

Gamut mapping is an ill-defined problem, and as such finding an optimal solution remains an open problem. Our gamut mapped results present differences with respect to the target BRDFs we are trying to represent. This is to be expected, since the target BRDFs lie outside the gamut, and therefore compromises need to be made when bringing them inside. These differences may be due to the inability of the inks to represent certain material properties (e.g., since there are no purely diffuse inks in our gamut, purely diffuse materials cannot be accurately represented), or to the optimization, since we cannot guarantee to find an optimal solution. Nevertheless, our approach yields better results in general than other state-of-the-art techniques.

Our attribute-based framework allows for versatility to achieve a variety goals, since different appearance properties can be preserved during the mapping process. As a consequence of this versatility, the particular choice of attributes may also have an influence on the final result, differing from the target BRDF. For example, *silvermetallicpaint* (Figure 10, top) benefits from the preservation of the *rough* attribute, since it is one of the main characteristics of the target BRDF, while preserving the *metallic-like* attribute instead would yield sharper reflections.

Finally, materials that lie far away from the gamut defined by the inks remain a challenging problem; in such cases our method may fail to faithfully reproduce the desired appearance, causing the resulting in-gamut BRDF to present visual differences with respect to the target BRDF. This behavior is similar in Pereira's work, suggesting that the limited gamut provided by the inks is the main cause for such differences in these cases. An example of this is depicted in Figure 12, showing also how the single-step state-of-the-art method fails. However, our result preserves the specular behavior better, thanks to our initial step in which we preserve the *metallic-like* attribute.

While currently users can choose to preserve different attributes with different weights, an interesting future line of work would be to conduct perceptual studies to analyze the influence of each attri-

bute in the perceived appearance, in order to automatically assign weights to each attribute during the optimization process. Further investigation could also be devoted to optimizing for just a region of the image that contains most of the appearance information, as opposed to the whole image.

Acknowledgements

We would like to thank the reviewers for their insightful comments, and Adrian Jarabo and Ibon Guillen for discussions. We also thank Sandra Malpica. This research has been partially funded by an ERC Consolidator Grant (project CHAMELEON), and the Spanish Ministry of Economy and Competitiveness (projects TIN2016-79710-P and TIN2016-78753-P). Tiancheng Sun was further supported by Tsinghua Spark Program. Ana Serrano was supported by an FPI grant from the Spanish Ministry of Economy and Competitiveness.

References

- [Deb98] DEBEVEC P.: Light probe image gallery. <http://www.pauldebevec.com/Probes/>, 1998. Last accessed May 2017. 4
- [DWP*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 62. 2
- [FDA03] FLEMING R. W., DROR R. O., ADELSON E. H.: Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3 (2003), 347–368. 2, 4, 5
- [FFGZ12] FORES A., FERWERDA J., GU J., ZHAO X.: Toward a perceptually based metric for BRDF modeling. In *20th Color and Imaging Conference* (2012), CIC'12, pp. 142–148. 4, 5
- [GGG*16] GUARNERA D., GUARNERA G., GHOSH A., DENK C., GLENCROSS M.: BRDF representation and acquisition. *Computer Graphics Forum* 35 (2016), 625–650. 2
- [HFM*10] HAŠAN M., FUCHS M., MATUSIK W., PFISTER H., RUSINKIEWICZ S.: Physical reproduction of materials with specified subsurface scattering. In *ACM Transactions on Graphics (TOG)* (2010), vol. 29, ACM, p. 61. 2
- [HFM16] HAVRAN V., FILIP J., MYŠKOWSKI K.: Perceptually motivated BRDF comparison using single image. *Computer Graphics Forum* 35, 4 (2016), 1–12. 2
- [LDPT13] LAN Y., DONG Y., PELLACINI F., TONG X.: Bi-scale appearance fabrication. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 145. 2
- [MAA01] MCCOOL M. D., ANG J., AHMAD A.: Homomorphic Factorization of BRDFs for High-performance Rendering. In *Proc. of SIGGRAPH '01* (2001), pp. 171–178. 3
- [MAG*09] MATUSIK W., AJDIN B., GU J., LAWRENCE J., LENSCH H. P. A., PELLACINI F., RUSINKIEWICZ S.: Printing spatially-varying reflectance. *ACM Transactions on Graphics (TOG)* 28, 5 (Dec. 2009), 128:1–128:9. 1, 2, 4
- [MMG11] MCNAMARA A., MANIA K., GUTIERREZ D.: Perception in graphics, visualization, virtual environments and animation. *SIGGRAPH Asia 2011 (Course)* (2011). 2
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics (TOG)* 22, 3 (jul 2003), 759–769. 1, 2, 4
- [NDM06] NGAN A., DURAND F., MATUSIK W.: Image-driven navigation of analytical BRDF models. In *Eurographics Symposium on Rendering (EGSR)* (2006), pp. 399–407. 1
- [NJR15] NIELSEN J. B., JENSEN H. W., RAMAMOORTHI R.: On optimal, minimal BRDF sampling for reflectance acquisition. *ACM Transactions on Graphics (TOG)* 34, 6 (November 2015), 186:1–186:11. 1, 2, 3
- [PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (2000), pp. 55–64. 2, 5
- [PR12] PEREIRA T., RUSINKIEWICZ S.: Gamut mapping spatially varying reflectance with an improved BRDF similarity metric. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 31, 4 (June 2012). 1, 2, 3, 4, 5, 7
- [SGM*16] SERRANO A., GUTIERREZ D., MYŠKOWSKI K., SEIDEL H.-P., MASIA B.: An intuitive control space for material appearance. *ACM Transactions on Graphics (TOG)* 35, 6 (nov 2016), 186:1–186:12. 1, 2, 3, 6, 7
- [VLD07] VANGORP P., LAURIJSSSEN J., DUTRÉ P.: The influence of shape on the perception of material reflectance. *ACM Transactions on Graphics (TOG)* 26, 3 (July 2007). 2
- [WAKB09] WILLS J., AGARWAL S., KRIEGMAN D., BELONGIE S.: Toward a perceptual space for gloss. *ACM Transactions on Graphics (TOG)* 28, 4 (Sept. 2009), 103:1–103:15. 2
- [WN99] WRIGHT S., NOCEDAL J.: Numerical optimization. *Springer Science* 35 (1999), 67–68. 3
- [WPMR09] WEYRICH T., PEERS P., MATUSIK W., RUSINKIEWICZ S.: Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (TOG)* 28, 3 (jul 2009), 32:1–32:6. 2