

SIGGRAPH: U: Attribute-preserving gamut mapping of measured BRDFs*

TIANCHENG SUN, University of California San Diego, Universidad de Zaragoza (I3A), and Tsinghua University
ANA SERRANO, DIEGO GUTIERREZ, and BELEN MASIA, Universidad de Zaragoza (I3A)

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. [2016], 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.

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 unconstrained problem without a unique solution. The goal of the problem is to find the BRDF which is *the most similar* to the target BRDF, while lying within the available gamut. To achieve this, we propose a two-step gamut mapping technique: In the first step, we leverage recent works on material acquisition [Nielsen et al. 2015] and editing [Serrano et al. 2016]. In these works, Nielsen et al. first built a 5D principal component (PC) space which serves as a basis for representing each BRDF; then, Serrano et al. learnt functionals mapping the space of principal components to higher level perceptual attributes defined for achromatic reflectance; these functionals define an intuitive control space for appearance. We use these mappings in PC space to bring the luminance channel L (L in Lab space) of the target BRDF into the 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 is still out of gamut (Figure 1). We thus complete the gamut mapping process with our second step, an image-based optimization, inspired by other recent works [Pereira and Rusinkiewicz 2012].

*This work was done during Tiancheng Sun’s internship in Universidad de Zaragoza, which was funded by Tsinghua University while he was an undergraduate there. Tiansheng Sun is currently a PhD student in UC San Diego.

2 RELATED WORK

Perceptual spaces for BRDFs. The derivation of perceptual spaces can be used for gamut mapping, in particular for the establishment of distances among BRDFs. Recently, Serrano et al. [2016] presented an intuitive space for BRDFs where a series of functionals are derived for connecting measured BRDF data with the perceptual ratings obtained from a large-scale user studies. In this work we use the functionals derived by Serrano et al. [2016] to map the original BRDF to the nearest material inside the gamut that preserves certain perceptual attributes.

Gamut mapping for BRDFs. Several approaches have been proposed to reproduce specific material appearances. Closer to our approach, some 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 Matusik et al. [2009], while Lan et al. [2013] calculate the L2 norm on the BRDF hemisphere for data fitting. Pereira and Rusinkiewicz [2012] improved this procedure by minimizing the difference between rendered images of the original and the final materials. 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 attribute of the original BRDF.

3 ATTRIBUTE-PRESERVING GAMUT MAPPING

Our goal is to take an out-of-gamut BRDF and bring it into a representable gamut, defined by the BRDFs of the individual color inks of a printer, while preserving a given perceptual attribute, such as its *brightness*. 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 the initial BRDF into gamut (in PC space); these isocontours represent the same value of a given perceptual attribute (please refer to [Serrano et al. 2016] for details). For visualization purposes we show a 2D slice of the original 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 [Pereira and Rusinkiewicz 2012]) 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 [2012]. We use the *alumina oxide* BRDF from the MERL database, and aim to preserve

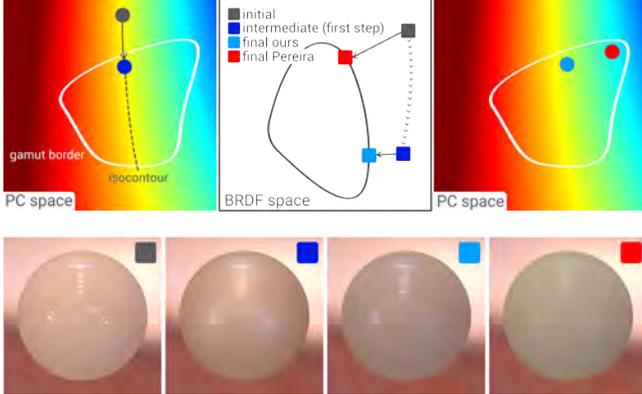


Fig. 1. Overview of our two-step gamut mapping method. **Top-left:** A 2D projection of the 5D PC space. The white line represents the border of the gamut, and same color-coded isocontours indicate the same value of a given perceptual attribute (following [Serrano et al. 2016]). Working on achromatic 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 [Pereira and Rusinkiewicz 2012]. **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 BRDF *alumina oxide* 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 [Pereira and Rusinkiewicz 2012]. Our result preserves highlights better, while exhibiting less color shift.

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.

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 the gamut in log-mapped PC space [Nielsen et al. 2015] following the path that maintains the same value v_A of a certain perceptual attribute A^1 . 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 here the same procedure. First, we apply a log-relative mapping to ρ_{ini} , which enables a good distribution of the available dynamic range [McCool et al. 2001; Nielsen et al. 2015; Serrano et al. 2016], and obtain the first five coefficients $\alpha_{ini} \in \mathbb{R}^5$ of the BRDF in the PCA basis (which provide general hints about material appearance [Nielsen et al. 2015]). Then, we apply a function $f_A: \mathbb{R}^5 \rightarrow \mathbb{R}$ that maps a BRDF in the aforementioned 5D PC space to its

¹For the sake of clarity, we explain our method for the simpler case of fixing just one attribute.

attribute value: $v_A = f_A(\alpha_{ini})$. The mapping function f_A is a radial basis function network (RBFN) based on Gaussian function, and its weighted is learnt from the perceptual ratings obtained from a large-scale user studies [Serrano et al. 2016].

For our gamut mapping, we formulate the objective function to maintain the initial attribute value $v_A = f_A(\alpha_{ini})$, so that the optimization moves along the corresponding isocontour of f_A as much as possible. 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 $Conv(\alpha_{inks})$ limited by the ink coefficients in the 5D PC space. The resulting optimization problem becomes:

$$\min_{\alpha} \|f_A(\alpha) - f_A(\alpha_{ini})\|^2 \quad \text{s.t.} \quad \alpha \in Conv(\alpha_{inks}) \quad (1)$$

In order to solve this optimization we use sequential quadratic programming (SQP) [Wright and Nocedal 1999] 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} . We can easily extend the procedure to preserve multiple attributes by employ a linear combination of them in the optimization.

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 $\mathbf{P}_{inks} = [\rho_1 \ \rho_2 \ \dots \ \rho_N]$ 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:

$$\rho_{fin} = [\rho_1 \ \rho_2 \ \dots \ \rho_N] \cdot \mathbf{w}, \quad (2)$$

where $\mathbf{w} \in \mathbb{R}^N$ consists of the coefficients for each ink.

In principle, we can minimize the distance between the intermediate BRDF ρ_{int} and the reproducible one ρ_{fin} 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 [Fores et al. 2012]. Therefore, we will instead minimize this distance in image space, as has been done in the past [Pereira and Rusinkiewicz 2012].

Usually, pixel values of a rendered image has a complicated relation with the BRDF values due to the indirect lighting. However, if we consider a purely opaque sphere, which is commonly used to visualize BRDFs, under the environment map, the light interacts only once on its surface before reaching the camera. Thus, the rendered image and the BRDF are linearly related as:

$$\mathbf{I} = \mathbf{R} \cdot \rho, \quad (3)$$

where \mathbf{R} is a matrix defining the linear mapping. Using Equation 3, we can change Equation 2 into:

$$\mathbf{I} = [\mathbf{I}_1 \ \mathbf{I}_2 \ \dots \ \mathbf{I}_N] \cdot \mathbf{w}, \quad (4)$$

where \mathbf{I}_i are the rendered images for each ink BRDF. Note that this sidesteps the need to explicitly compute \mathbf{R} ; moreover, these rendered

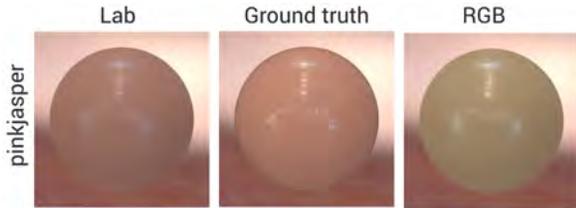


Fig. 2. Gamut mapped results for the *pinkjasper* MERL BRDF (middle) optimizing in Lab (left) and RGB (right). Image-based optimization in Lab space better preserves chromaticity. The set of inks that define the gamut can be seen in Figure 3.

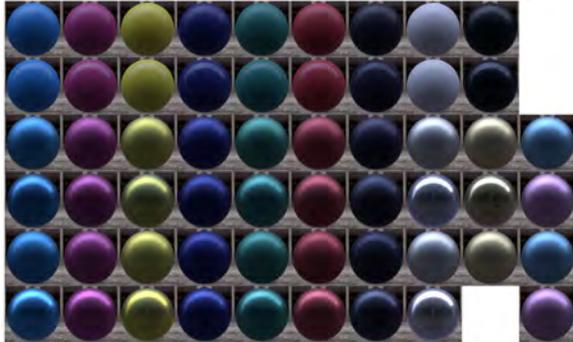


Fig. 3. All the BRDFs from the gamut provided by Matusik et al. [2009], which are measured from real inks.

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} :

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

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

where \mathbf{I}_{int} is the image obtained with the BRDF computed in the first step. We choose the distance d in image space to be the L_2 norm under Lab color space, since it better preserves the color of the original BRDF, as shown in Figure 2.

4 RESULTS

In our results, we use the BRDF gamut from Matusik et al. [2009]. 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 4.1) we use the *Eucalyptus Grove* map², given that these illuminations facilitate material perception [Fleming et al. 2003], 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.

²Both environment maps are from the Light Probe Image Gallery [Debevec 1998].

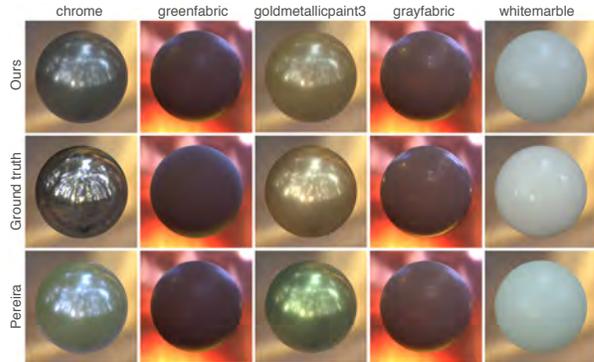


Fig. 4. Comparison of our results and the state-of-the-art method by Pereira [2012] using different illuminations (*Eucalyptus* and *Grace*). In general, our method minimizes color shifts (*chrome* and *goldmetallicpaint3*), while better preserving highlights and specular behavior (*grayplastic* and *whitemarble*). For very diffuse materials (*greenfabric*) neither method succeeds due to the specular nature of the inks used Figure 3.

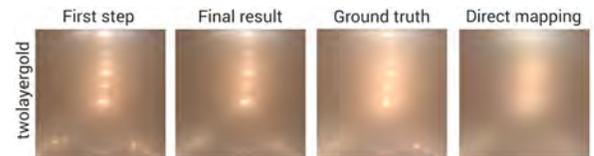


Fig. 5. Our two-step, attribute-preserving gamut mapping, compared to a single-step optimization. Note how our method helps to preserve the specular highlights and overall appearance better.

We compare our results with those from Pereira [2012] on 94 homogeneous materials from the MERL database [Matusik et al. 2003]. Some representative results rendered with different environment maps are shown in Figure 4.

4.1 Validation

Two-step validation. Figure 5 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 [2012]) can also be seen in Figure 1.

Objective validation. To provide an objective comparison to the state of the art [Pereira and Rusinkiewicz 2012], we use the cube root cosine weighted RMS metric, which has been reported to perform better than RMS for BRDFs [Fores et al. 2012]. This metric is described as:

$$E = \sqrt{\frac{\sum_n ((\rho_{\text{fin}} \cos \theta_i - \rho_{\text{ini}} \cos \theta_i)^2)^{1/3}}{n}} \quad (5)$$

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

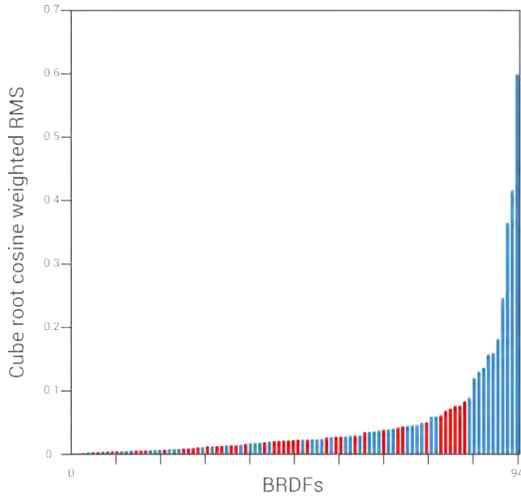


Fig. 6. Difference between the error of the state-of-the-art image-based optimization [Pereira and Rusinkiewicz 2012] and that of our method; the error is computed as the cube root cosine weighted RMS [Fores et al. 2012]. Blue indicates better results with our method.

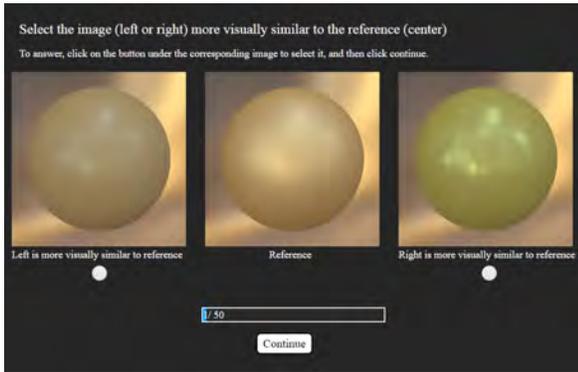


Fig. 7. Screenshot of our user study. The reference is presented in the middle, with the two options at the sides in random order.

higher). Although we do not outperform the single-step method of Pereira [2012] for all BRDFs in the database, we do in a majority of cases.

User study. We have also 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 [Fleming et al. 2003; Pellacini et al. 2000; Pereira and Rusinkiewicz 2012] we use a sphere to depict the materials. 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 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.

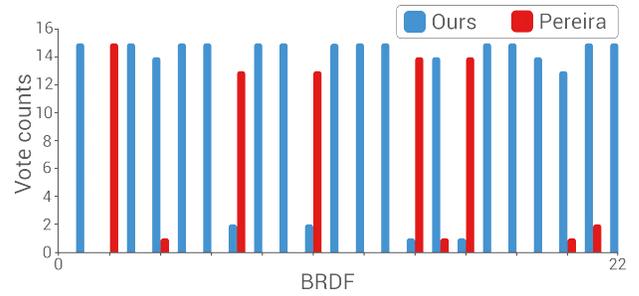


Fig. 8. Vote counts indicating preference for the BRDFs mapped with our method (blue bars) and Pereira’s method (red bars) for the BRDFs with statistically significant differences. In them, our result was preferred 17 out of 22 times with high agreement between users.

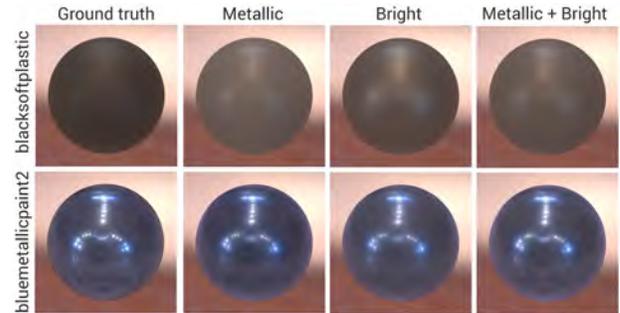


Fig. 9. From left to right: Original BRDF and corresponding gamut mapping results computed by preserving, during the first step of our method, only the *metallic-like* attribute, only the *bright* attribute, and both attributes. Optimizing over the *metallic-like* isocontour yields more accurate reflections, 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.

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.

4.2 Preserved perceptual attributes

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. 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. Here, we further 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. When optimizing to preserve only the *metallic-like* attribute, the resulting BRDF preserves the specular behavior better,

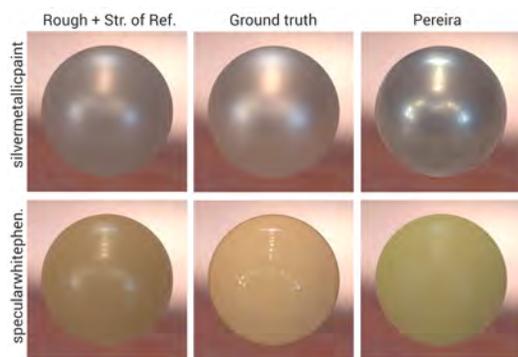


Fig. 10. Results computed preserving both the *rough* and *strength of reflections* attributes in the first step of the method, and comparison to the state of the art [Pereira and Rusinkiewicz 2012]. This particular combination of attributes aims to better preserve the appearance of the reflections.

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. Figure 10 shows additional results preserving 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.

5 DISCUSSION AND CONCLUSION

In this paper, we have proposed a 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 perceptual ratings 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. Our attribute-based framework allows for versatility to achieve a variety of 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.

Gamut mapping is an ill-defined problem, and finding an optimal solution still 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 a global optimal. Nevertheless, our approach yields better results in general than other state-of-the-art techniques.

Materials that lie far away from the gamut remain a challenging problem; in such cases our method may fail to faithfully reproduce



Fig. 11. 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.

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 11, 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 attribute 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.

REFERENCES

- Paul Debevec. 1998. Light Probe Image Gallery. <http://www.pauldebevec.com/Probes/>. (1998). Last accessed May 2017.
- Roland W. Fleming, Ron O. Dror, and Edward H. Adelson. 2003. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3 (2003), 347–368.
- A. Fores, J. Ferwerda, J. Gu, and X. Zhao. 2012. Toward a Perceptually based Metric for BRDF Modeling. In *20th Color and Imaging Conference (CIC'12)*. 142–148.
- Yanxiang Lan, Yue Dong, Fabio Pellacini, and Xin Tong. 2013. Bi-scale appearance fabrication. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 145.
- Wojciech Matusik, Boris Ajdin, Jinwei Gu, Jason Lawrence, Hendrik P. A. Lensch, Fabio Pellacini, and Szymon Rusinkiewicz. 2009. Printing Spatially-varying Reflectance. *ACM Transactions on Graphics (TOG)* 28, 5, Article 128 (Dec. 2009), 9 pages.
- Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. 2003. A Data-Driven Reflectance Model. *ACM Transactions on Graphics (TOG)* 22, 3 (jul 2003), 759–769.
- Michael D. McCool, Jason Ang, and Anis Ahmad. 2001. Homomorphic Factorization of BRDFs for High-performance Rendering. In *Proc. of SIGGRAPH '01*. 171–178.
- Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. On Optimal, Minimal BRDF Sampling for Reflectance Acquisition. *ACM Transactions on Graphics (TOG)* 34, 6 (November 2015), 186:1–186:11.
- Fabio Pellacini, James A Ferwerda, and Donald P. Greenberg. 2000. Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. 55–64.
- Thiago Pereira and Szymon Rusinkiewicz. 2012. Gamut Mapping Spatially Varying Reflectance with an Improved BRDF Similarity Metric. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 31, 4 (June 2012).
- Ana Serrano, Diego Gutierrez, Karol Myszkowski, Hans-Peter Seidel, and Belen Masia. 2016. An intuitive control space for material appearance. *ACM Transactions on Graphics (TOG)* 35, 6, Article 186 (nov 2016), 186:1–186:12 pages.
- Stephen Wright and Jorge Nocedal. 1999. Numerical optimization. *Springer Science* 35 (1999), 67–68.