

# Gamut Mapping Spatially Varying Reflectance with an Improved BRDF Similarity Metric

Thiago Pereira and Szymon Rusinkiewicz

Princeton University

---

## Abstract

Recent spatially varying reflectance (svBRDF) printing systems can reproduce an input document as a combination of matte, glossy and metallic inks. Due to the limited number of inks, this reproduction process incurs some distortion. In this work, we present an svBRDF gamut mapping algorithm that minimizes distortions in the angular and spatial domains. To preserve a material's perceived variation with lighting and view, we introduce an improved BRDF similarity metric that builds on both experimental results on reflectance perception and on the statistics of natural lighting environments. Our experiments show better preservation of object color and highlights, as validated quantitatively as well as through a perceptual study. As for the spatial domain, we show how to adapt traditional color gamut mapping methods to svBRDFs. Our solution takes into account the contrast between regions, achieving better preservation of textures and edges.


Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Reflectance, gamut mapping

---

## 1. Introduction

While most real-world materials exhibit a variety of appearances ranging from diffuse to glossy to metallic, traditional printing methods can only reproduce grayscale or color images. Motivated by advances in printing technologies, recent work has shown how to create physical reproductions of materials with angular dependent effects. These include opaque materials [MAG\*09], but also translucent ones [DWP\*10, HFM\*10]. In these works, the authors note that practical applications require goal-based printing pipelines, i.e. the user should only specify the output appearance instead of how much ink or material is to be used in each place. For instance, Matusik et al. [MAG\*09] describe a reflectance printing system in which the user inputs a *spatially varying bidirectional reflectance distribution function* (svBRDF) [NRH\*92], and the system finds the proper combination of inks of different reflectances through halftoning.

However, one challenge in traditional printing is also present in goal-based reflectance printing: gamut mapping [Mor08]. Any printer has only a few different inks available, and it is not possible to achieve exact reproduction of input documents. These distortions may be in the angular dimension (e.g. highlights are not broad enough) or in the spatial dimension (e.g. edges have reduced contrast). In this work, we address the problem of svBRDF gamut mapping: finding



**Figure 1:** We address the problem of mapping a BRDF to a constrained gamut, such that it is close according to some similarity metric. Our metric (right) achieves better reproductions of the target material (left) when seen inside natural lighting environments compared to Matusik et al. (middle left) [MAG\*09] and Pellacini et al. (middle right) [PL07].

the best possible approximation to an input svBRDF in the reproducible set of the printer, i.e. the printer's BRDF gamut.

While a large literature exists for image gamut mapping [Mor08, CIE04, LHM11, KSES05], these algorithms cannot handle the angular effects of svBRDFs. To minimize angular domain distortion, it is possible to use a BRDF similarity metric [PFG00, NDM06, PL07]. Matusik et al. [MAG\*09] used a metric that optimizes reproduction for point light visualization, which does not necessarily lead to good reproductions under natural environments. The distortion can be significant, specially for metals and specular materials. Inspired by the statistics of lighting environments [DLAW01], we propose a metric based on a new synthetic environment

that correlates well with natural environments. We show that using this metric reduces perceptual distortion (Figure 1).

In addition, previous work has ignored the spatial arrangement of BRDFs, which may lead to loss of contrast in edges and texture. Finding a method that preserves the contrast between BRDFs and scales to the size of svBRDF datasets is a challenge. We have adapted recent gamut mapping approaches [KSES05, LHM11] and show how they perform on svBRDF datasets.

Our method has applications beyond reproduction. For instance, it could fit parametric models to captured svBRDFs. In addition, our metric could be used to achieve more perceptually accurate svBRDF decompositions [LBAD\*06] and interactive edits [PL07].

Our main contributions are:


- An improved perceptual BRDF similarity metric based on a new synthetic lighting environment that correlates with natural environments (Section 4). We validate it through gamut mapping experiments and a perceptual user study.
- An adaptation of an image gamut mapping algorithm for svBRDFs. Our solution builds on our metric for clustering and optimal projections, but also takes the spatial arrangement of the BRDFs into account (Section 5). Our experiments show this method better preserves textures and edges (Section 6).

## 2. Related Work

**BRDF similarity metric:** An important part of an svBRDF gamut mapping system is its BRDF metric. Pellacini et al. [PFG00] proposed a perceptually uniform reparametrization of the low-dimensional monochrome Ward BRDF space based on a psychophysical study. Generalizing this work to the high-dimensional space of real BRDFs, however, would be impractical. This has led to the use of algorithms, rather than explicit perceptual measurements, for determining BRDF similarity.

One approach are metrics that have an analytical expression, such as the  $L^2$ -norm and the metric of Lawrence and Pellacini [PL07], even though neither is perceptually inspired. Ngan et al. proposed a perceptual image-based metric [NDM06] inspired by studies showing that humans are best at judging reflectance when seen under natural environments [FDA01]. Therefore, they propose to compare BRDFs by instead comparing environment-mapped rendered images of spheres made of the materials (although [VLD07] argues that more complex shapes are better for human perception).

However, Ngan et al. do not recommend a specific environment to be used as metric. Our experiments show that this choice can make a big difference. Matusik et al. [MAG\*09] used a single point light as environment in their svBRDF reproduction system. While it leads to a simple metric, they claim that the point light does not correlate well with natural



**Figure 2:** a) Reflectance inks from [MAG\*09] b) Rendered sphere in our environment and its coordinate systems.

environments when testing on large datasets of BRDFs. Our environment BRDF metric also follows this image-based approach, leveraging its corresponding perceptual results, but in addition we introduce a synthetic environment for comparing BRDFs that is consistent with the statistics of natural environments [DLAW01].

While a BRDF metric should be consistent with high-level reflectance perception, it should also build on low-level perception of color. For this, we use the CIELAB metric [Fai05], even though other metrics could be used.

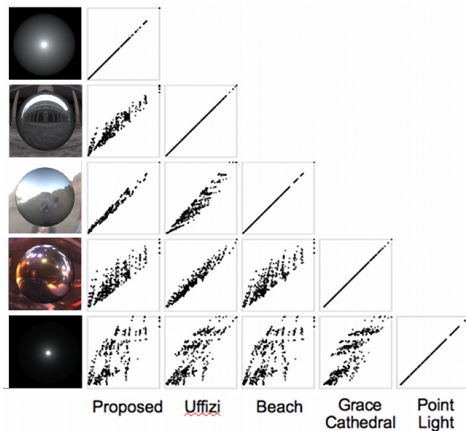
**svBRDF gamut mapping:** When mapping svBRDFs, it is not enough to maximize similarity per pixel. Instead of a reproduction that approximates the original in an absolute sense, we should rather aim at *relative* reproduction of the svBRDF, which aims at preserving edges and textures.

A good survey on color gamut mapping can be found in [Mor08]. However, it's unclear how to extend many color-specific concepts that are central to these algorithms, such as hue preservation, luminance remapping and black point compensation. More recent developments, which we extend to svBRDFs, are spatial gamut mapping algorithms. They allow the same color to be mapped differently depending on its spatial position by using signal processing [ZS07, FGR07] or optimization methods [KSES05, LHM11].

In addition to spatial domain similarity, Guthe et al. [GMSK09] also consider the angular domain and develop a metric for Bidirectional Texture Functions. However, their metric predicts just-noticeable differences, while we focus on larger differences. The first work that focuses on large-scale changes in the angular domain for svBRDF gamut mapping was Matusik et al. [MAG\*09]. In this work, the authors preserve spatial details by mapping material bases in the svBRDF convex hull and preserving combination weights. In their work, a BRDF is always mapped the same, independently of its position. In contrast, we allow it to map differently in order to preserve contrast.

Hersch et al. [HCE03] describe a reproduction system with color and metallic inks and Stollnitz et al. [SOS98] with multiple color inks. However, both works mainly focus on predicting the appearance of a combination of inks while we focus on minimizing perceptual BRDF distortion.

Other works gamut map materials with subsurface scattering appearance [DWP\*10, HFM\*10]. Due to the more limited availability of scattering inks, these works focus on reproducing the achromatic characteristics of the materials.



**Figure 3:** Scatterplots showing correlations between BRDF similarity under different environments. Each plot considers a pair of environments, with each point corresponding to a pair of BRDFs. Its  $x$  and  $y$  coordinates are the distances between the materials as measured in both environments (so that points closer to the diagonal indicate better agreement). Our synthetic  $1/f$  environment is shown in the leftmost column: it has high correlation with the natural environments.

### 3. Reproduction Framework

Matusik et al. [MAG\*09] present a reflectance printing system that receives an input svBRDF and maps it to the printer gamut. We follow many of their assumptions about the printing process. First, they capture the BRDFs of a variety of ink stacks (e.g. a stack of cyan, yellow and silver foil ink would yield a metallic yellow ink). In this work, whenever we talk about inks, we refer to these stacked composites. They also discuss how through halftoning a printer can generate any convex combination of its basis inks. In short, in this work, we assume the printer gamut to be any convex combination of their composite basis inks (Figure 2). As a source gamut we use the MERL database [MPBM03], containing a representative set of 100 BRDFs, and we consider several svBRDFs captured in previous work [LBAD\*06].

We represent all these BRDFs as  $\rho(\theta_h)$  curves: sampled one-dimensional functions of the half-angle  $\theta_h$ , defined to be the angle between the surface normal and the bisector of view and light direction [Rus98]. This bisector gives the direction a microfacet would need to be oriented for perfect mirror reflection between light and observer. Since our input svBRDF maps and the output inks all have isotropic reflectance, this representation captures the main visual features of these BRDFs, namely the color and shape of the highlights. Its major disadvantage is that it does not model phenomena such as retro-reflection and grazing-angle effects. We find that previously proposed BRDF metrics performed well for some  $\theta_h$  curves, but also disagreed widely with human perception for others. For this reason, we present an improved BRDF metric in the next section.

### 4. BRDF Similarity Metric

In this section, we present an improved environment-based BRDF metric. Our metric builds on the idea of comparing two BRDFs by comparing rendered images of objects having the two BRDFs under natural environments [NDM06]. While it has been shown that humans can better perceive reflectance when seeing complex shapes [VLD07], we follow Ngan et al. [NDM06] and use simply the sphere shape, which leads to a simpler metric. After rendering, we convert each pixel's color to the CIELAB color space and compare the resulting images, taking an  $L^p$  difference pixel-wise. By doing so, we model human color perception more accurately.

However, we find the result of this metric to depend on the environment chosen. We designed a synthetic environment (subsection 4.1) that predicts well many natural environments because it is inspired by experimental analysis of natural environments. We propose its use as a reference for comparing BRDFs. We also show (subsection 4.2) how to find an analytical expression for our metric. In addition, we discuss two important features of our metric: use of CIELAB and choice of  $L^p$ -norm (subsection 4.3). Finally, we validate our metric by applying it to the gamut mapping problem and through perceptual user studies (subsections 4.4 and 4.5).

#### 4.1. Synthetic Environment

In this subsection, we describe our proposed synthetic environment, but first we present a comparison of multiple environment metrics for the gamut mapping problem. We compare the metrics induced by several reference environment maps (beach, Grace, Uffizi, kitchen, St Peters) on three sets of BRDFs: the APLS printer inks, the MERL database and a set of synthetic Ward BRDFs [War92] with  $k_s$  and  $k_d$  ranging from 0 to 1 ( $k_s + k_d \leq 1$ ) and roughness ranging up to 0.25. For each set, we compute pairwise similarities between all pairs of BRDFs under one environment, then compare the distances to those obtained using a different environment map. Ideally, we would like the BRDF similarity values to be consistent across different environments, in other words to have perfect correlation. Indeed, on the inks dataset, the correlation is quite high. However, for the MERL and Ward datasets there are substantial differences: though there is still a clear correlation between the results (Figure 3) using different environment maps, relative distances could vary by well over a factor of 2. It is interesting to notice that the Uffizi environment was most "typical" in the sense of agreeing best on average with other environments. Grace was most atypical. In conclusion, the environment-based metric is dependent on the environment that is used.

Let us consider next the metrics induced by two synthetic environments: a point light source at the camera and our new proposed environment. The latter also has a singularity at the pole, but has a heavy tail (Figure 3), with energy distributed as  $1/\tan(\theta/2)$ . It is motivated by the observation that real-world environments tend to have total energy per frequency

varying roughly as  $1/f$  [DLAW01] (sum of  $2l + 1$  spherical harmonics coefficients of energy  $1/l^2$ ). This is achieved with an environment having energy distributed roughly as  $1/\theta$ . We actually choose  $1/\tan(\theta/2)$ , so that the function goes to zero as theta goes to  $\pi$ . We refer to this as the  $1/f$  environment and its corresponding metric as the **1/f metric**.

In our experiments with the nine Debevec environments, we found significant greater variance in energy distribution than is acknowledged in previous work [FDA01, DLAW01]. While we also find the mean energy per spherical harmonic coefficient to be  $1/l^2$ , the exponents range between 1.5, 2 (St. Peters, Galileo), 2.1 ( $1/f$  environment) up to 3 and 4 (Uffizi, Beach). Any environment that we pick is a compromise and we must evaluate how it agrees with others as a metric.

Comparing the BRDF metrics induced by point and  $1/f$  to those of real environments, we find that they all have high correlation on inks. On Ward and MERL, the point light source is worse than any real environment, while the  $1/f$  environment is comparable (Figure 3). In particular, it has high agreement with beach. This is easy to understand, since beach has essentially a major source of light, i.e. the sun, but also a radial falloff as light scatters off the sky. The conclusion is that the analytic  $1/f$  environment does about as well as any real environment, plus it is radially symmetric (leading to faster evaluation).

For gamut mapping applications, we can replace any metric by its composition with an increasing function. Therefore, one could argue that correlation between metrics is not an appropriate measure of similarity between metrics. For this reason, we additionally ran all our correlations experiments using Spearman's correlation. This is a correlation function that is invariant to composition with increasing functions. While some details do change, our conclusions are essentially the same.

One important advantage of the  $1/f$  environment is its symmetry. In the next section, we show how it can be used to find an analytical expression for our metric.

#### 4.2. Expression in Half-Angle Coordinates

In this section, we exploit the symmetry of our proposed environment and the symmetry of our chosen BRDF representation of  $\theta_h$  curves to find an analytical expression for our metric, which can be used as an alternative to actually rendering the images. This expression is easier to plug in optimization methods. The process of computing the  $1/f$  metric on  $\theta_h$  curves is complex, since it requires a full spherical convolution. Let us now derive our metric equation and show how to precompute these convolution weights.

We start from the illumination equation to calculate the image we would obtain by rendering a sphere of constant BRDF described by a  $\theta_h$  curve under any radially symmetric environment. Given two different BRDFs, we can integrate

the  $L^p$  difference of their images. We assume that both the viewer and the environment's symmetry center are in the up direction (Figure 2). Parameterizing the sphere with  $\theta_o, \phi_o$ , we can see that radiance arriving at the eye from the sphere only depends on  $\theta_o$ . Using the area element of the projected sphere in this parametrization and simplifying:

$$\frac{d(\rho_1, \rho_2)}{2\pi} = \int_0^{\pi/2} d_c(I_1(\theta_o), I_2(\theta_o))^p \sin \theta_o \cos \theta_o d\theta_o, \quad (1)$$

where  $d_c$  is any color space metric. At this point, we look into the spherical convolution that results in the rendered image  $I(\theta_o)$ . This means we now integrate over incident light directions  $\omega_i$  for a fixed  $\theta_o$  value. We denote the viewer direction in the incident hemisphere's coordinate system  $\omega_o$ . The outgoing light in this direction can be calculated by the following integral  $I(\theta_o) = \int f(\omega_i, \omega_o) E(\omega_i) \cos \theta_i d\omega_i$ . Integrating instead in the  $\theta_h, \phi_h$  coordinates [Rus98] where the BRDF is a function of a single variable  $\rho(\theta_h)$ :

$$I(\theta_o) = \int \rho(\theta_h) E(\omega_i) \cos \theta_i \left| \frac{\partial \omega_i}{\partial \theta_h} \times \frac{\partial \omega_i}{\partial \phi_h} \right| d\phi_h d\theta_h. \quad (2)$$

We can precalculate all that does not depend on the BRDF:  $I(\theta_o) = \int A(\theta_o, \theta_h) \rho(\theta_h) d\theta_h$ . We now make the expression of the function  $A$  more explicit. For our setup, we know that  $\omega_o$  is in the z-direction (Figure 2) and we rewrite the environment  $E(\omega_i) = E(\angle(\omega_i, \omega_o)) = E(2\theta_d)$ . We also know that  $\cos \theta_i = \omega_i^z$  and from the definition of  $h$ :  $\cos \theta_i = 2 \cos \theta_d \cos \theta_h - \cos \theta_o$ . We also need the  $\theta_h, \phi_h$  area element [Ren50]. In addition, by solving  $\cos \theta_i > 0$ , we find the range of integration of  $\phi_h$  to be  $[-\cos^{-1}K, \cos^{-1}K]$ , where  $K = -\cot \theta_o \cot 2\theta_h$ . Substituting, we obtain  $A(\theta_o, \theta_h) =$

$$4 \sin \theta_h \int E(2\theta_d) (2 \cos \theta_d \cos \theta_h - \cos \theta_o) \cos \theta_d d\phi_h.$$


We believe that it is not possible to obtain a closed-form expression for this integral for our environment. For this reason, we discretize and precalculate  $A(\theta_o, \theta_h)$  numerically. In conclusion, these weights let us compute  $I(\theta_o)$  as a matrix-vector multiplication, which is easier to discretize and optimize. This formulation, lets us evaluate our metric in 3 ms in MATLAB. Our equivalent implementation rendering using a modern graphics card and BRDF importance sampling, but no symmetry, takes 120 ms. This speed up becomes crucial when processing svBRDFs.

#### 4.3. Color and Image Comparison

We use the CIELAB color metric (D50 illuminant) since it builds on perceptual experiments on human color perception. While the RGB Euclidean metric could be used, our experiments find this solution to be inappropriate. In many cases, the RGB metric leads to incorrect hue, which can happen for both diffuse and glossy materials (Figure 4-d).

The integral in Equation 1 of our metric is essentially a simple image comparison metric. While we could have used more complex image metrics [WBSS04], the simpler choice





**Figure 4:** These comparisons display the target material (left) and the result with our metric (right). Our BRDF metric is built on four main ideas. First, that BRDFs should be compared through their rendered images in lighting environments. Image a) shows a reproduction using  $L^2$ -norm in  $\theta_h$ -space with CIELAB color comparison. Our result preserves the hue. Second, that color comparisons should be performed using a perceptual color metric. Image d) middle shows the results of our metric using RGB space. Third, that to preserve highlights, the integration over angle should be performed as an  $L^4$  norm. Middle of images b) and c) show our metric using  $L^2$  instead, which results in blurred highlights. Finally, that BRDFs should be compared in the  $1/f$  environment. For comparison, images e) and f) middle present the result using the point light metric [MAG\*09].

of an  $L^p$  norm was favored. We experimentally chose  $p = 4$ . Compared to  $L^2$ , the  $L^4$  norm gives less weight to small pixel errors and more weight to large errors. This is consistent with image perception by humans, since we are very tolerant of small changes in the mean intensity of an image.

Experimenting with gamut mapping, we find the highlights to be particularly sensitive to changes in  $p$ . Figure 4-b,c shows a comparison of reproductions using  $p = 2, 4$ . The highlights are sharper with  $L^4$ , while the overall colors barely changed. While highlights continue to improve for higher values of  $p$ , this leads to deviations in diffuse color.

#### 4.4. BRDF Mapping Results

Having fully described the proposed BRDF metric, we now present comparison to previously proposed metrics using gamut mapping experiments. The images shown next are the result of solving a gamut mapping optimization problem. The mapping of a BRDF  $y$  onto the gamut of a set of inks  $W$  using the metric  $d$  can be written as:

$$\begin{aligned} \min_x \quad & d(Wx, y) \\ \text{s.t.} \quad & x \geq 0, \sum x_i = 1. \end{aligned}$$


where  $Wx$  is a convex combination of the columns of  $W$ , i.e.,  $Wx$  is any ink in the gamut defined by the convex hull of our basis inks. This is a non-linear optimization problem when using the CIELAB metric, but is a simpler quadratic problem in the RGB case. While most previous methods have not used the CIELAB metric for color comparison, we chose to implement them with CIELAB for a more fair comparison. We next show a comparison between ours and the  $L^2_{\theta_h}$ , point and cosine metrics. We discuss how they compare under different environments and gamuts.

We begin by showing gamut mapping results using the  $L^2_{\theta_h}$  norm:  $\int_0^{\pi/2} d_c(\rho_1(\theta_h), \rho_2(\theta_h))^2 d\theta_h$  where  $d_c$  is the CIELAB metric. As can be seen in Figure 4-a, even when using the CIELAB metric, this simple  $L^2_{\theta_h}$  metric fails to reproduce hue in many cases. The main drawback of this metric is that it does not compare BRDFs under any kind of lighting environment. The extreme mistakes above can be avoided by using an environment metric such as the point light. However, as we show below, our proposed metric can still achieve better results compared to the **point light metric** (Figure 5). The reason behind these results, as we saw earlier, is the low correlation between natural environments and the point light. In Figure 5, we show the exact same materials illuminated under a point light. The point metric indeed results in more similar materials when seen under point lighting.

We show many different cases where our metric achieves improved results. In Figure 4, our main advantage is the correct overall color. In other cases, as in Figure 5, the improvement is in the highlight color and sharpness: notice how our reproduction shows the reflection of the buildings.

We also compare our metric to the weighted **cosine metric** [PL07]. We implemented it in RGB space because it led to a simpler quadratic optimization problem. Therefore, it is only fair to compare to our metric for gray materials. Our experiments show that this metric behaves similar to the point CIELAB metric. They both give too much weight to the highlights and tend to miss the overall color (Figure 1).

Overall, for our mapping experiments on the MERL database, our metric leads to less perceptual distortion when the materials are seen under natural environments like beach, kitchen and Uffizi. As for Grace and St Peters, the point light




**Figure 5:** We achieve better reproductions under natural environments like Uffizi and beach, while the point metric is better under point and Grace. However, a detailed analysis (dashed lines) of Grace shows that in regions dominated by an area light, our proposed metric is better. Within each set: target (left), point (middle) and our metric (right). From top to bottom: visualization under Uffizi, point and Grace. We clip high intensity pixels.




**Figure 6:** Restricted to a diffuse gamut, our metric (right) preserved the color, as opposed to the point metric (middle).

metric results in less perceptual distortion on average (Figure 5). This can be understood because these environments are composed of a large collection of distinct point lights. However, in regions where the environment is an area light, our metric more closely matches the original (Figure 5).

These results show that our metric can achieve good reproductions of the MERL materials on the inks gamut. It appears that most noticeable artifacts still left are a result of gamut limitations, e.g. the inks gamut does not include very diffuse materials, very glossy materials or very dark materials. We push these limitations further by removing the basis BRDFs that used silver or gold foil inks in their composition (columns 4,6,10 and 12 in Figure 2), which are necessary to approximate some metallic materials. We show results when mapping aluminium, but we observe similar behavior with many metallic materials. In Figure 7, we can see how our metric prefers a reconstruction with a broad highlight, which is certainly far from the target due to gamut limitations. The other metrics prefer a very dark reconstruction in order to have a sharper highlight. In these extreme trade-offs it is not so clear what is desirable, but, in our opinion, our reproduction was more faithful. Even only using diffuse inks, our metric can still create a reasonable reproduction (Figure 6).



**Figure 7:** If we do not employ foil inks, we cannot reproduce aluminium (left). The  $L_{0_h}^2$  and point light metrics (middle) preferred dark results in order to have a sharper highlight. Our metric (right) preferred a broader highlight.



**Figure 8:** User interface used in our perceptual user studies.

Our comparisons to previous metrics under different environments and gamuts have shown the improved perceptual quality of our metric. To further validate it, we ran two perceptual user studies.

#### 4.5. Perceptual Studies

We designed two perceptual studies to compare the performance of our solution to two other metrics in the gamut mapping problem. For each study, we used a different selection of materials from the MERL database. Our selected gamut was the same set of BRDF inks discussed in previous sections. We chose three representative environments based on our previous correlations analysis: beach, Uffizi and Grace.

Both studies use an interface (Figure 8) where the target material is presented in the middle and mappings with two different metrics are presented at the sides. The subjects are asked to select which of the mappings is more similar to the middle one. In addition, we also give them the option “Equally similar”. By clicking and holding the mouse button, the user can temporarily observe the left and right images swapped. This overlaying makes it faster to observe some of the appearance differences. We also randomize left and right. This kind of perceptual study retains the disadvantage of the study of Pellacini et al. [PFG00], i.e. only images are compared instead of the actual physical material.

We recruited 16 subjects between 21 and 37 years old. There were 11 males and 5 females. Some of the subjects participated in both studies. The average session took 25 minutes with 9 seconds per question. All the sessions were performed on the same calibrated display. Each sphere displayed occupied roughly 3.5 degrees of visual angle.

In our first experiment, we compare our metric to the point

Ours preferred over	Uffizi and Beach	Grace
point CIELAB	73%	41%
cosine RGB	74%	67%

**Table 1:** Frequency with which the  $1/f$  metric is preferred. Statistical significance  $p < 0.05$  for  $\pm 5\%$  confidence.

light CIELAB metric. A share of the MERL database lies inside the inks gamut. Since, for these cases, both metrics yield essentially the same results, we decided to focus on out of gamut materials. For this purpose, for each environment, we calculate CIELAB pixel differences between the two mappings. For beach/Uffizi/Grace, we discard material environment pairs where the average pixel color difference is less than 1/1/1 JND (just noticeable difference) and maximum less than 3/3/6 JND. We believe these thresholds are conservative because 50% of the materials remain and subjects find that many materials are still similar. A total of 153 questions are equally distributed among environments.

We observe that in this setting only 14% of the responses were ‘Equally similar’. This means that our subjects found significant enough differences between the two metrics to justify a selection. We also observe that our subject population was rather consistent among themselves. On average, only 14.5% of the responses were the opposite of the majority (e.g. majority prefers metric A and response prefers B). In addition, we found this majority is usually significant: on average 75% of the subjects agree on a choice. The consistency of this population leads us to believe that the observed preferences are generalizable to a larger population.

For the beach and Uffizi environments, we find that in most cases where subjects have a preference they prefer our metric compared to the point metric. It was chosen 73%  $\pm$  2% ( $p < 0.05$ ). For the Grace environment, they prefer our metric in only 41%  $\pm$  3.5% ( $p < 0.05$ ). These results are consistent with our discussion in the previous section.

In our second experiment, we compare our metric to an RGB implementation of the cosine metric [PL07]. Since we are comparing against an RGB implementation, we restricted our study to only approximately monochrome materials in the MERL database, not necessarily out of gamut. This study consisted of 35 materials for a total of 105 questions equally distributed among the three environments. Since many materials are in gamut, we observed a higher rate of ‘Equally similar’ responses 31%. We again observed a consistent population, only 8% of the responses are the opposite of the majority opinion. Even though the target materials are monochrome, the mappings with the cosine metric often result in some chroma. This led to users preferring our metric in 74%  $\pm$  4% ( $p < 0.05$ ) of the questions for Uffizi and beach and 67%  $\pm$  6% ( $p < 0.05$ ) for Grace.

All these results demonstrate a frequent preference for our metric. This preference is also considerable, which is attested by the non-forced choice nature of our study.

## 5. svBRDF Mapping

After finding improved BRDF mappings, we now consider multiple BRDFs per image, i.e. an svBRDF. Possibly the simplest approach to svBRDF mapping is to consider each pixel’s BRDF in isolation and map it to the closest in gamut BRDF. This approach is known as **clipping**. Another approach is the **convex compression** solution presented in Matusik et al. [MAG\*09]. In their work, they represent each BRDF in the document by a convex combination of basis BRDFs. They choose the basis near the convex hull of the source gamut, so that, when these are mapped, all the other BRDFs are compressed inside the destination gamut as well.

Both of these methods have advantages and drawbacks. Clipping has the property that it does not change materials that are already in gamut. However, it can lose spatial details in regions of the svBRDFs where all pixels map to the same in-gamut color. Convex compression can introduce large changes even to in-gamut colors, which leads to a loss of global contrast. In cases where multiple basis clip to the same point, compression also leads to loss of details.

To overcome these limitations, we have adapted two algorithms [KSES05, LHM11] from the spatial gamut mapping literature. Both techniques use optimization to preserve point-wise BRDF similarity and their spatial differences. They can be written in the following general form:

$$\begin{aligned} \min_{x_p} \quad & \sum_{p \in V} d(Wx_p, y_p) + \alpha \sum_{(u,v) \in E} d(Wx_u - Wx_v, y_u - y_v) \\ \text{s.t.} \quad & x_p \geq 0, \sum x_p^i = 1, \forall p \in V. \end{aligned}$$

where  $x_p^i$  are all the ink weights associated with vertex  $p$ ,  $y_p$  is the target BRDF at vertex  $p$  and  $W$  is the gamut matrix as described in the previous section. Our objective function is non-linear because similarity is measured using our environment metric  $d$ . In fact, clipping can also be written in this general form by setting  $\alpha = 0$ .

The major difference between these two algorithms is how to define the sets of vertices  $V$  and edges  $E$ . The first is the **gradient-based** method [KSES05]. It aims at preserving spatial gradients (i.e. difference between neighboring pixels). In other words, choose  $V$  to be the set of pixels and  $E$  to be the edges in all four neighborhoods. While this solution does manage to achieve a balance between clipping and compression algorithms, often having the advantages of both, it also suffers from two major drawbacks. It creates halo artifacts around strong edges (Figure 9). Even though halo could be improved with sparse gradient norms, this solution is also very slow, on the order of an hour for a 30 by 30 image. Compared to gradient reconstruction methods, this problem is harder because it is non-linear and constrained, which renders common speed-up techniques inapplicable.

A solution to these two problems is preserving the contrast between *regions* instead of pixels. This fixes halo because it focuses the optimization at preserving significant

contrasts, as opposed to all pixel differences. In addition, there is a major speed-up since the number of variables becomes much smaller, proportional to the number of regions.

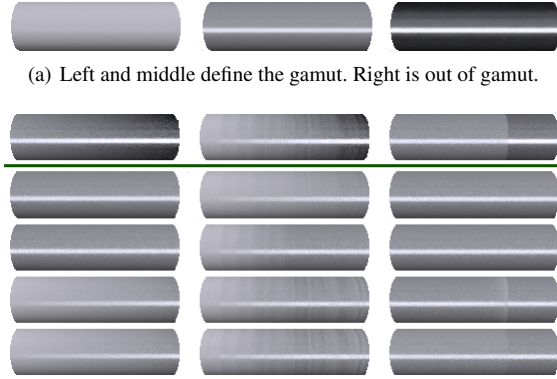
This method is similar to the optimization based algorithms of Lau et al [LHM11]. It consists of four parts: clustering, optimization, interpolation and clipping. First, it starts by clustering the pixels into regions defined by BRDF and spatial distance with k-means. Clustering is performed in the perceptual space implied by our metric appended by the two spatial coordinates. In other words, we use the  $I(\theta_o)$  curves in CIELAB space weighted by  $\sqrt{\cos\theta_o \sin\theta_o}$  as our feature vectors. Second, the same optimization above is applied by taking V to be the cluster BRDF means and E to be neighboring clusters. This procedure preserves the differences between adjacent regions. Third, we add back the details. The simplest way would be to add the displacement between the mapped and source cluster material to all the other pixels in this cluster, but this may lead to artifacts at clustering boundaries. To avoid this, we follow the approach of Lau et al. [LHM11], in which the authors interpolate the displacement vectors  $W_{x_p} - y_p$  based on the inverse distance to each cluster center. Finally, to ensure that the image is in gamut, we run a final clipping step. This whole process has the advantage of good global contrast as a result of the optimization but also good preservation of details. We will refer to this method as the **cluster-based** solution.

In conclusion, clipping and convex compression fail to preserve contrast and edges. The gradient and cluster-based methods do a good job at preserving contrast, but only the cluster-based scales to the size of svBRDF datasets. Next section shows results that support these claims.

## 6. svBRDF Results

To compare these methods, we present three simple but challenging use cases. We use the same gamut in all cases: the convex combinations of the diffuse and specular inks shown in Figure 9-a. The metallic ink also shown in this figure is out of gamut and is mapped to the specular ink by our metric. Each svBRDF is a horizontal gradient of two of these three materials (Figure 9). We choose to visualize them wrapped around cylinders because they are isometric to the document plane, but let us see multiple orientations in one single image. As a result, in a single cylinder any left to right variation is due to material change while top to bottom variations are caused by illumination. The number of clusters and  $\alpha$  are inputs to the algorithm (values used are shown in captions).

In the first column, the input cylinder (top) is completely out of gamut: a gradient between specular and metallic. Both convex mapping and clipping lose all spatial variation. Gradient and cluster preservation preserve the spatial variation by using some diffuse ink on the left. The second input cylinder is only half out of gamut: a gradient with noise added to the combination weights. Convex compression preserves



(a) Left and middle define the gamut. Right is out of gamut.

**Figure 9:** Comparison of four svBRDF gamut mapping algorithms. Each column shows the mappings of a different cylinder. The target (above green line) is mapped using convex compression (first below green line), clipping (second), gradient preservation (third) and the cluster-based solution (bottom). Only the cluster-based solution works in all cases.

the gradient variations, but compresses the details. Clipping perfectly reproduces the left half of the cylinder, but loses all variations in the right half. Again, gradient and cluster preservation achieve good reproductions. Our third case is a simple edge between out of gamut materials. Convex compression and clipping lose the edge. Gradient preservation preserves the edge but leads to halo artifacts. The cluster-based solution preserves the edge without halo. Overall, the cluster-based method gives good results in all cases.


We also ran experiments on the svBRDFs from Lawrence et al. [LBAD\*06]. We found the full set of 57 inks to result in very good reproductions for this svBRDF dataset. To create more challenging cases, many of the experiments presented include gamuts with a reduced number of inks, gamuts with darker inks or svBRDFs where we retained the spatial variations but replaced the basis materials.

In Figure 11, all inks were scaled by 90% and we also mixed 25% diffuse white into all pixels to move the image more out of gamut. In the middle, we show how clipping can lead to total loss of edges. On the right, the cluster-based method can recover the contrast and make the edges visible again at the cost of worse matching of specularly.

As another example, we replaced the basis materials for the wheel. Our target appearance is made of a convex combination of brass, mapped brass and perfect black (Figure 13). Combinations of mapped brass and brass are projected to the same material. This leads to a strong loss of texture when using the simple clipping method. For both the wheel and the season greetings card, the convex compression method of Matusik et al. [MAG\*09] would lead to similar results to clipping because these are cases where multiple basis materials are mapped to the same place.

The target in Figure 10 is a combination of copper and





**Figure 10:** The clipping output is closer to the target (left) but loses global contrast. The cluster-based result (right) is darker in the red region enhancing the contrast. Visualized with our environment (top triangle) and point light (bottom). Generated with  $\alpha = 50$  and 25 clusters.




**Figure 11:** Clipping (middle) can lead to total loss of edges from the target (left). In this case, the cluster method can recover the contrast and make the edges visible (right) at the cost of worse matching of specularities. Visualized with our environment (top triangle) and point light (bottom). Generated with  $\alpha = 250$  and 10 clusters.

brass, both out of gamut. The clipping output is closer to the target but loses the contrast between red and yellow regions. The cluster-based method leads to a darker reproduction, but recovers the contrast. It preserves the specularities of the input, which is only possible by using the gold and silver foil inks. This explains the reproduction's lack of red.

As a final example, we studied the behavior of our metric with progressively restricted gamuts (Figure 12) using the clipping method. The target dove (first column) is well reproduced by the full gamut (second column). In the third column, we removed the metallic inks. This led the system to use the foil inks, which are more specular than desired. For the fourth column, we removed both metallic and foil inks. The result still shows a very narrow highlight from the inks that include a finish layer. The fifth column displays the svBRDF clipped using only the diffuse inks. Some specularities can still be seen since these are not perfectly diffuse.

## 7. Implementation Details

**Discrete metric:** Discretizing our metric is straightforward. All functions of  $\theta_h$  and  $\theta_o$  become vectors by sampling the angles uniformly and transforming the integrals into summations. Equation 2 defines a linear mapping between the space of  $\theta_h$  curves and the space of  $\theta_o$  curves, which we write in vector notation as:  $I = A\rho$ . Each of its entries is precalculated by numerical integration. Our metric can be computed by



**Figure 12:** On the left we show the target svBRDF, which is well reproduced by the inks (second column). In the third column, we show the result using a gamut without metallic inks. In the fourth column, we further remove the foil inks. Finally, we show the projection on a diffuse gamut. Top row shows the dove visualized with the point light environment, bottom row with a  $1/f$  environment.



**Figure 13:** The cluster-based method (right) preserves texture, while clipping destroys it (middle). Generated with  $\alpha = 1.25$  and 25 clusters.

applying this matrix to the reflectance vectors, transforming the result to CIELAB and replacing the image space integral by a weighted summation.

**Optimization:** We implemented the BRDF gamut mapping process using an interior point method available through MATLAB's `fmincon` function. Because of the CIELAB non-linearity, our metric is a non-convex function. However, in practice, we did not observe local minima. We initialize the algorithm with a uniform combination of all inks. All the svBRDF algorithms were also implemented using the interior-point method in MATLAB. All are initialized with uniform inks. Most of the implementation details of the cluster-based method can be found in Lau et al. [LHM11], but there are two major differences that we find necessary to make this solution practical for svBRDFs. First, in their work, they formulate this problem with a quadratic objective function subject to non-linear gamut constraints. Instead, we chose to formulate the non-linear objective function but linearly constrained problem. Second, even though projecting a single BRDF to a gamut of 57 inks takes only 2 seconds, we find it computationally prohibitive to simply project all pixels in our final svBRDF clipping step. Instead we quantize the input svBRDF with a large number of clusters (on average 400) using k-means. We only project the centers of these clusters. We visually inspected all images to make sure this quantization step is introducing negligible distortion.

## 8. Conclusion and Future Work

We presented a new gamut mapping algorithm for svBRDFs. We have shown how a synthetic environment can lead to a simple metric, but still agree with most natural environments resulting in perceptually accurate reproductions, including material color and highlights. We have also adapted some existing image gamut mapping methods to the svBRDF context. We show how the cluster-based solution leads to good preservation of textures and edges, avoids halo artifacts and scales well to size of the svBRDFs.

One limitation of our approach is that we restricted the BRDF to depend only on  $\theta_h$ . Therefore, we cannot represent retro-reflection, grazing-angle and anisotropic appearance. An interesting question is how to approximate anisotropic by isotropic BRDFs. In addition, our method is limited to svBRDFs, in which all interaction happens at the surface. We would like to extend our metric to scattering materials. Another future direction is extending our methods to svBRDFs on height maps or even 3D surfaces, allowing its application in 3D printing.

## Acknowledgements

Much of this work originated in conversations with Wojciech Matusik and Fabio Pellacini, and we thank them for many helpful suggestions. We also thank Wojciech Matusik, Boris Ajdin, Jason Lawrence, and Paul Debevec for the BRDF, svBRDF, and environment-map datasets used in this paper. Finally, we thank the National Science Foundation (grants CCF-1012147 and CCF-1027962) and the Intel Science and Technology Center for Visual Computing for their support.

## References

- [CIE04] CIE: Guidelines for the evaluation of gamut mapping algorithms. In *Technical Report CIE 156:2004* (2004). 1
- [DLAW01] DROR R., LEUNG T., ADELSON E., WILLSKY A.: Statistics of real-world illumination. In *Computer Vision and Pattern Recognition*. (2001). 1, 2, 4
- [DWP\*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.* 29 (July 2010), 62:1–62:10. 1, 2
- [Fai05] FAIRCHILD M. D.: Color appearance models. John Wiley and Sons. 2
- [FDA01] FLEMING R., DROR R. O., ADELSON E. H.: How do humans determine reflectance properties under unknown illumination. In *In Proceedings of CVPR Workshop on Identifying Objects Across Variations in Lighting: Psychophysics and Computation* (2001), pp. 347–368. 2, 4
- [FGR07] FARUP I., GATTA C., RIZZI A.: A multiscale framework for spatial gamut mapping. *Image Processing, IEEE Transactions on* 16, 10 (oct. 2007), 2423–2435. 2
- [GMSK09] GUTHE M., MÜLLER G., SCHNEIDER M., KLEIN R.: Btf-cielab: A perceptual difference measure for quality assessment and compression of btfs. *Computer Graphics Forum* 28, 1 (Feb. 2009), 101–113. 2
- [HCE03] HERSCH R. D., COLLAUD F., EMMEL P.: Reproducing color images with embedded metallic patterns. *ACM Trans. Graph.* (2003). 2
- [HFM\*10] HAŠAN M., FUCHS M., MATUSIK W., PFISTER H., RUSINKIEWICZ S.: Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph.* (2010). 1, 2
- [KSES05] KIMMEL R., SHAKED D., ELAD M., SOBEL I.: Space-dependent color gamut mapping: a variational approach. *Image Processing, IEEE Transactions on* (2005). 1, 2, 7
- [LBAD\*06] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHY R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.* (2006). 2, 3, 8
- [LHM11] LAU C., HEIDRICH W., MANTIUK R.: Cluster-based color space optimizations. In *Proceedings of International Conference on Computer Vision (ICCV 2011)* (2011). 1, 2, 7, 8, 9
- [MAG\*09] MATUSIK W., AJDIN B., GU J., LAWRENCE J., LENSCH H. P. A., PELLACINI F., RUSINKIEWICZ S.: Printing spatially-varying reflectance. *SIGGRAPH Asia*. 1, 2, 3, 5, 7, 8
- [Mor08] MOROVIC J.: Color gamut mapping. Wiley. 1, 2
- [MPBM03] MATUSIK W., PFISTER H., BR M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics* 22 (2003), 759–769. 3
- [NDM06] NGAN A., DURAND F., MATUSIK W.: Image-driven navigation of analytical brdf models. In *Eurographics Symposium on Rendering* (2006). 1, 2, 3
- [NRH\*92] NICODEMUS F. E., RICHMOND J. C., HSIA J. J., GINSBERG I. W., LIMPERIS T.: Radiometry, geometrical considerations and nomenclature for reflectance. 1
- [PFG00] PELLACINI F., FERWERDA J. A., GREENBERG D. P.: Toward a psychophysically-based light reflection model for image synthesis. *SIGGRAPH*. 1, 2, 6
- [PL07] PELLACINI F., LAWRENCE J.: Appwand: editing measured materials using appearance-driven optimization. *ACM Trans. Graph.* 26 (July 2007). 1, 2, 5, 7
- [Ren50] RENSE W. A.: Polarization studies of light diffusely reflected from ground and etched glass surfaces. *J. Opt. Soc. Am.* 40, 1 (Jan 1950), 55–56. 4
- [Rus98] RUSINKIEWICZ S. M.: A new change of variables for efficient brdf representation. In *In Eurographics Workshop on Rendering* (1998), pp. 11–22. 3, 4
- [SOS98] STOLLNITZ E. J., OSTROMOUKHOV V., SALESIN D. H.: Reproducing color images using custom inks. *SIGGRAPH*. 2
- [VLD07] VANGORP P., LAURIJSSSEN J., DUTRÉ P.: The influence of shape on the perception of material reflectance. *SIGGRAPH '07, ACM*. 2, 3
- [War92] WARD G. J.: Measuring and modeling anisotropic reflection. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1992), *SIGGRAPH '92, ACM*, pp. 265–272. 3
- [WBSS04] WANG Z., BOVIK A., SHEIKH H., SIMONCELLI E.: Image quality assessment: from error visibility to structural similarity. *Image Processing, IEEE Transactions on* (2004). 4
- [ZS07] ZOLLIKER P., SIMON K.: Retaining local image information in gamut mapping algorithms. *Image Processing, IEEE Transactions on* 16, 3 (march 2007), 664–672. 2