

# Printing Spatially-Varying Reflectance

Wojciech Matusik<sup>1</sup> Boris Ajdin<sup>1,2</sup> Jinwei Gu<sup>1,3</sup> Jason Lawrence<sup>4</sup>  
Hendrik P. A. Lensch<sup>5</sup> Fabio Pellacini<sup>1,6</sup> Szymon Rusinkiewicz<sup>1,7</sup>

<sup>1</sup>Adobe Systems, Inc. <sup>2</sup>MPI Informatik <sup>3</sup>Columbia University <sup>4</sup>University of Virginia <sup>5</sup>Ulm University <sup>6</sup>Dartmouth College <sup>7</sup>Princeton University



**Figure 1:** *Top left:* Our system represents a target document as a spatially-varying Bidirectional Reflectance Distribution Function (svBRDF). We have developed BRDF gamut-mapping and halftoning algorithms that approximate this svBRDF with a set of printer inks. *Bottom left:* We visualize how the document will appear when printed and observed under specified viewpoint and lighting. *Right:* Printed sample.

## Abstract

Although real-world surfaces can exhibit significant variation in materials—glossy, diffuse, metallic, etc.—printers are usually used to reproduce color or gray-scale images. We propose a complete system that uses appropriate inks and foils to print documents with a variety of material properties. Given a set of inks with known Bidirectional Reflectance Distribution Functions (BRDFs), our system automatically finds the optimal linear combinations to approximate the BRDFs of the target documents. Novel gamut-mapping algorithms preserve the relative glossiness between different BRDFs, and halftoning is used to produce patterns to be sent to the printer. We demonstrate the effectiveness of this approach with printed samples of a number of measured spatially-varying BRDFs.

## 1 Introduction


Digital publishing has revolutionized the way we create, edit, and view documents. More importantly, the availability of inexpensive yet high-quality home printers has greatly simplified the process of producing physical copies of digital documents. Over the last thirty years, digital publishing has grown to become a multi-billion dollar industry, and it has affected the life of artists, professionals, and home users alike.

However, one major limitation of current digital documents, as well as many printing and display devices, is that they represent images using only gray-scale or color values. This is a drawback since the appearance of real-world surfaces cannot be represented faithfully by color or gray values alone. For example, real-world surfaces can exhibit different amounts of speculariry: they range from perfectly matte (diffuse) to glossy to mirror-reflective. They can also appear metallic, with colored specular highlights. Therefore, there exists a large gap between the way we represent digital documents and the way their physical equivalents look in reality.

This stands in contrast to the steady progress in the range of appearance properties that some printers can output. For example, even now it is possible to use matte, glossy, or metallic inks and apply highly specular overcoats to printed surfaces, to achieve a large variety of appearances. In order to take full advantage of these capabilities, however, it is necessary to use a richer representation for digital documents. It is also required to redefine the process of converting this representation to a physical copy.

We propose to represent a digital document using a spatially-varying reflectance function instead of gray-scale or color values. In particular, we use a spatially-varying Bidirectional Reflectance Distribution Function (svBRDF), which we will write as  $f_r(x, \omega_i, \omega_o)$ . This function encodes, for each point  $x$  on the document's surface, the ratio of the outgoing light (radiance) in a direction  $\omega_o$  to the incident light (irradiance) from a direction  $\omega_i$ . Knowing the svBRDF, we can accurately predict how the document will appear, from any desired viewpoint and under any lighting configuration. The goal of our publishing process can therefore be stated as starting with the reflectances (BRDFs) of some set of available inks, and reproducing the goal svBRDF as faithfully as possible with a spatial pattern of those inks.

Our key observation is that, through spatial halftoning of inks, we can reproduce any BRDF within the convex hull of our ink library. Thus, we do not restrict the artist to using individual inks as “spot colors,” but rather provide a full spectrum of materials. For example, Figure 2 shows a halftoned ramp between two inks of different specularities, oriented to reflect a checkerboard target. As




**Figure 2:** *Top:* A printed ramp between “silver foil” and “metallic silver” inks, set up to reflect a checkerboard target. The perceived specularly of the surface varies smoothly from left to right. *Bottom:* Closeups of the halftoning patterns used near the left and center of the pattern.

can be seen, the effect of dithering is to give the impression of new BRDFs that are linear combinations of the inks — this is a direct consequence of the linearity of light transport.

We believe that the availability of reflectance functions as first-class citizens in document editing programs will improve users’ access to the richness available using a variety of inks in today’s printing technologies. Moreover, by allowing the user to think in terms of *materials* instead of individual *inks*, we gain two benefits. First, we can leverage the fact that combinations of inks may be automatically employed to broaden the range of available appearance. More importantly, we free the user from the constraints and tedium of worrying about spot colors and explicit separation into ink layers, thus enabling greater freedom and creativity.

Our process (Figure 3) begins with the target svBRDF, as well as measurements of the printer BRDFs: layered combinations of the available inks on a given substrate. Although we could simply map each target BRDF to the closest printable BRDF (as discussed in Section 3.2), this would lead to undesirable clamping artifacts if target BRDFs fall outside the gamut of the available inks. Instead, we draw inspiration from the *gamut-mapping* process performed by all (color) printers (Section 4.1). We find the “most extreme” BRDFs present in the document, such that the document is representable as a convex linear combination of them, then map only those BRDFs to the printer gamut. All other points are represented as the corresponding linear combinations of the remapped BRDFs. Finally, we transform continuous per-pixel ink weights into discrete halftone patterns, at the same or higher resolution. Our halftoning algorithm (Section 4.2) is inspired by those used for colors, but is adapted for BRDFs and incorporates physical dot-size constraints of our printer. The final patterns can be previewed (under arbitrary view and lighting), and then are sent to the printer.

We validate our reflectance printing process using both measured and synthetically generated spatially-varying reflectance functions as input. We show simulations and the corresponding printouts for our pipeline, demonstrating the effects of the BRDF gamut-mapping and halftoning stages. Our target is a desktop thermal printer (ALPS MD5500), and our printer BRDF library contains 57 different combinations of inks, foils, and finishes, all captured using an image-based reflectance measurement process (Section 5.2).



**Figure 3:** *Flowchart of our system.* We start with a measured or synthetically generated svBRDF, which can be modified by applying user edits. Next, the svBRDF is projected onto a space spanned by the BRDFs of available inks. Gamut mapping is employed in order to avoid clamping artifacts. Then, halftone patterns are generated for each of the ink BRDFs in order to approximate their continuous weights. Finally, the svBRDF can be previewed under arbitrary lighting and from arbitrary viewing directions, or it can be sent to a printer to generate a physical copy.

## 2 Related Work

The prior work for this paper spans two different research areas that have traditionally had little overlap. First, we review work related to printing, focusing on research in custom printing and halftoning algorithms. Second, we describe prior work in appearance representation and perception.

**Custom Printing:** Stollnitz et al. [1998] investigated the problem of printing color images with multiple custom inks, proposing novel physical models, gamut-mapping and halftoning algorithms that consider overprinting, trapping, dot gain, and the interreflection of light between ink layers. Hersch et al. [2007] created gamut-mapping and halftoning methods for printing images with custom fluorescent inks: the printed images are invisible under daylight and indoor lighting, but become visible when illuminated with UV light. Most relevant to our work is the method of Hersch et al. [2003], who combined a metallic ink with standard inks to print color images with embedded metallic patterns. Such printed images have locations that exhibit a particular amount of specular reflection. In contrast, our goal is to print images that can continuously change appearance for different views, according to the desired material properties.

**Halftoning:** Since it was invented in the early 1850s, physical halftoning has been extensively used in the printing industry for generating grayscale and color images with a reduced number of inks, while appearing similar to the original image. Digital halftoning methods do not have the restrictions of physical screens, and a variety of algorithms have been developed; detailed surveys can be found in [Ulichney 1987; Kang 1999]. Recent developments in halftoning algorithms have focused on designing optimal templates that adapt to the target images. Ostromoukhov [2001] developed novel error diffusion templates that are optimal in their blue-noise properties. Pang et al. [2008] proposed a global optimization-based halftoning algorithm that minimizes both the color and structural difference between the halftoned image and the original. Inspired by these algorithms for grayscale or color images, this work develops halftoning algorithms for printing spatially-varying BRDFs.

**Appearance Representations:** Bidirectional Reflectance Distribution Functions are typically represented using low-parameter analytical formulas [Cook and Torrance 1982; Ward 1992; Lafortune et al. 1997], which can be fitted to reflectance measurements of real surfaces [Ngan et al. 2005]. Measured BRDFs can also be simply tabulated [Matusik et al. 2003], or approximated with factored representations (e.g., using sums of products of lower-dimensional functions) [Kautz and McCool 1999; McCool et al. 2001; Lawrence et al. 2006]. There is also a significant body of work on measuring and representing spatially-varying reflectance functions. These reflectance functions can be tabulated [Dana et al. 1999; Marschner et al. 1999], or represented using per-pixel fits to analytical re-



**Figure 4:** We use a factored representation for svBRDFs, in which each pixel is represented as a linear combination (visualized using the weight maps in the bottom row) of basis BRDFs (visualized in the top row as spheres rendered in the “Uffizi” environment).

flectance models [McAllister 2002; Gardner et al. 2003] or linear combinations of basis BRDFs [Lensch et al. 2003; Goldman et al. 2005; Lawrence et al. 2006].

**Material Perception:** Understanding the perception of material properties by the human visual system is essential for evaluating the visual differences between the desired and printed surface, and is thus an important component of our printing pipeline. However, this is still an area of active research. Pellacini et al. [2000] proposed a mapping, based on user-study data, from the Ward BRDF model’s parameters to a perceptually-uniform “gloss” space. Fleming et al. [2001] showed that humans can recognize reflectance under natural illumination with high accuracy from a single image. Inspired by this work, Ngan et al. [2006] proposed an image-based metric for evaluating BRDFs and developed a system that allows a user to navigate BRDF models for designing material properties. We leverage these perceptual results in our printing algorithms.

### 3 Document Representation

Our fundamental representation for the target document is an svBRDF. In this section we discuss our factored svBRDF representation, as well as our curve-based BRDF representation and how we assign a distance metric to it.

#### 3.1 Factored svBRDF Representation

Despite their flexibility, svBRDFs have the drawback of requiring significant storage, when represented uncompressed. To address this, it has become common to restrict the possible space of representable BRDFs at each pixel. This may be done by using analytic BRDF models, with only a small number of parameters stored at each pixel, or through the use of factored models, in which each pixel is a linear combination of basis materials.

We choose to use a factored svBRDF model (Figure 4), since it provides a combination of flexibility, accuracy, and space efficiency [Lawrence et al. 2006]. In particular, we have observed that many of our inks are not easily representable using analytic models such as those of Phong [1975] or Ward [1992]. The ability to represent those inks explicitly as bases in the output of our gamut-mapping stage is vital, thus strengthening our motivation for using a factored svBRDF.

In generating a factored svBRDF from measured data, we constrain the weight maps at each pixel to form a non-negative partition of unity (i.e., each weight must be between 0 and 1, and their sum must be 1). Although there are typically many factorizations satisfying this constraint, we are especially interested in those having basis BRDFs that lie on or near the convex hull of the BRDFs in the document: this is important for our gamut-mapping approach (Section 4.1). We have observed that in many instances the Sparse Alternating Constrained Least-Squares (SACLS) approach intro-

duced by Lawrence et al. [2006] yields such results, and in those cases we simply use the output of SACLS. In other instances, we adopt a modified  $k$ -means clustering in which, following each  $k$ -means iteration, we remove from consideration those points that are well-represented as linear combinations of the cluster centers. In this way, the cluster centers naturally migrate towards the convex hull of the BRDFs, since they are no longer attracted by BRDFs *within* the convex hull. Of the results shown in this paper, the “Season’s Greetings” and “Dove” datasets were produced using SACLS, while the “Oak Leaf” and “Rusty Flange” datasets used the convex  $k$ -means algorithm.

#### 3.2 Curve-Based BRDF Representation

For each of the target and printer BRDFs, we also desire a compact and flexible representation. Because of physical constraints of the printing process, we restrict the representation to isotropic materials. Moreover, we recognize that specular-highlight shape is the most salient feature of the different inks available to us, and so we seek a representation that can accurately describe the color and shape of highlights, while sacrificing the representability of phenomena such as retroreflection and grazing-angle effects.

We choose to represent each BRDF as  $f_r = f(\theta_h)$ : a sampled curve of values along the halfway angle  $\theta_h$ , which is defined as the angle between the surface normal and the bisector of the view and light directions [Rusinkiewicz 1998]. We find that this representation is sufficient to capture most of the perceived differences between printer and target BRDFs, while remaining compact (we typically use between 32 and 512 samples in  $\theta_h$ , depending on the dataset). Such a representation has become widely used in the community in recent years, since it directly captures the distribution of microfacet normals for isotropic microgeometry-based BRDFs [Torrance and Sparrow 1967; Ashikhmin et al. 2000; Ngan et al. 2005; Han et al. 2007]. Note that our processing pipeline is not specific to  $\theta_h$  curves: we could instead accommodate any other linear (e.g., tabulated or factored) BRDF representation.

**BRDF Distance Metric:** In order to perform operations such as minimum-error projection of BRDFs, it is necessary to have a distance metric. We adopt the approach of Ngan et al. [2006], who observed that simple Euclidean distance between BRDF functions corresponds poorly to material perception. Instead, that work proposed comparing *environment-mapped renderings* of spheres of the given material, with a nonlinearity (per-channel cube root) applied to the images to approximate color difference perception.


We have compared variants of this metric, using several different environment maps, on our database of printer BRDFs. We find good agreement among the distances computed using different environments, and additionally find good agreement with purely single-directional illumination. It is important to note that this observation no longer holds when considering a greater variety of BRDFs: there is non-negligible disagreement among different environments, and generally poor agreement with directional lighting. Nevertheless, given our target application, we choose to employ the directional-light metric, which permits a significant simplification in the case of  $\theta_h$  curves (see Appendix).

### 4 Software Printing Pipeline

As shown in Figure 3, once we have a target svBRDF (either measured or specified by the user) we must project it onto the space spanned by the printer BRDFs. In most cases, this requires a gamut-mapping step, performed on the entire dataset.

#### 4.1 Projection onto Printer BRDFs and Gamut Mapping

In some cases, each of our target BRDFs (i.e., each spatial location of our svBRDF) can be expressed exactly as a linear combination of the printer BRDFs. In many other cases, such a linear combination



**Figure 5:** A simple example of a BRDF gamut-mapping algorithm in 1D. The input document contains 4 target BRDFs that increase in specularity (first column). The printer BRDFs (second column) cover a much smaller range of specularity. Clamping occurs when directly projecting target BRDFs onto printer BRDFs (third column). Gamut mapping of BRDFs (fourth column) preserves relative differences between remapped BRDFs.

is not exact but nevertheless represents a good approximation. For these datasets, at each pixel we directly optimize for linear blending weights that reproduce the target BRDFs:

$$\{\alpha_1 \dots \alpha_K\} = \arg \min d\left(f(\omega_i, \omega_o), \sum_{k=1}^K \alpha_k f_k(\omega_i, \omega_o)\right) \quad (1)$$

subject to

$$0 \leq \alpha_k \leq 1 \text{ and } \sum_{k=1}^K \alpha_k = 1,$$

where  $f(\omega_i, \omega_o)$  is the target BRDF,  $f_k(\omega_i, \omega_o)$  are the printer BRDFs,  $\alpha_k$  are the blending weights we must compute, and  $d$  is a distance function (discussed above). In order to simplify the above notation, we have assumed that:


- Each of the  $f_k(\omega_i, \omega_o)$  is the BRDF of a particular sequence of inks overlaid on top of the substrate layer.
- The substrate alone is also one of the  $K$  “printer BRDFs.”
- Except as accounted for above, there is no overprinting. That is, each dot is either printed with one of the “printer BRDFs” (which may itself correspond to a stack of several actual inks) or not printed at all.

These assumptions hold for many modern consumer printing technologies.

The reason we formulate (1) as an optimization is to allow for an arbitrary distance function. Of course, when the target BRDFs are exactly expressible as linear combinations of the inks, we can always drive the objective function to zero. However, when we must *approximate* the target BRDF, we wish to use an objective function that minimizes the perceptual difference between the target and the approximation.

When the range of target BRDFs is larger than that of the printer BRDFs, we must perform gamut mapping. As with gamut mapping of colors, the goal is to compress the range while preserving the perceptual similarity of remapped BRDFs: two pixels in the target that have similar materials should continue to have similar materials when re-mapped, while relative differences between materials should be preserved.

For example, suppose we have four BRDFs we want to print, with relative glossiness of 100%, 80%, 50%, and 10%, as shown in Figure 5. Furthermore, suppose we have only two inks, with



**Figure 6:** We remap basis materials to the closest ones representable as a convex linear combination of the printer inks, according to the metric in (2). Note that this metric must make a tradeoff between preservation of diffuse color and preservation of the specular peak. Here, the closest match was achieved with a slight shift in hue, particularly visible in the component at right.

glossiness of 10% and 50%. Simply mapping each target BRDF to the closest ink would cause clamping, as illustrated in the third column. That is, the BRDFs in the second, third, and fourth rows are no longer distinct, but have been mapped to the same printer BRDF. In contrast, a correct gamut mapping will map the maximum-glossiness target BRDF to the maximum-glossiness ink, and will remap the other target BRDFs to preserve their *relative* level of gloss.

Our approach to gamut mapping is simply to map each *basis* material of our factored target svBRDF to the closest BRDF representable as a linear combination of printer BRDFs. That is, instead of performing the optimization of equation 1 independently for each pixel, we only perform it once for each target basis BRDF. The spatial mixing maps are *retained* without change, meaning that intermediate BRDFs will again be mapped to intermediate linear combinations on output. This approach will only be successful if the svBRDF basis contains “extreme” materials — those that cannot themselves be represented as linear combinations of others. It is for this reason that we rely on a basis produced by either the SACLs or convex-hull  $k$ -means algorithms, as discussed in Section 3.1.

We note that this approach necessarily maps each target BRDF into a convex linear combination of the printer BRDFs. This is because each target BRDF is, by construction, a convex combination of the target basis BRDFs, which in turn are expressed as convex combinations of the printer BRDFs. We are therefore relying on the fact that the composition of two convex linear combinations is itself a convex combination. Thus, the remainder of our pipeline can assume that all BRDFs will fall within the space spanned by the printer BRDFs. We also note that this approach is computationally efficient (seconds per dataset), since it requires performing the constrained nonlinear optimization only for a small number of basis BRDFs. In contrast, projecting each BRDF independently would result in computation times of hours or days.

An example of our mapping is presented in Figure 6. Note that this stage will not preserve the original BRDFs exactly: some compromise is always necessary. For this reason, it is essential that the user be able to preview the results of gamut mapping, which we pro-

vide in the form of renderings of spheres under environment-map lighting (as shown here) or full svBRDF renderings under multiple light directions (Figure 1).

**Discussion:** The gamut-mapping approach we use may be thought of as an analogue of traditional color-mapping algorithms that employ a *perceptual* rendering intent, with the exception that we adapt to the materials present in a target svBRDF. This has the advantage of maximizing the range of BRDFs that are represented in a printout, but with the disadvantage that the same BRDF may be remapped differently in different target svBRDFs. An alternative would have been to employ a “soft clamp” for target BRDFs near or outside the convex hull of printer BRDFs: a strategy inspired by color-mapping approaches that preserve hue and value, while gradually clamping saturation. The per-pixel strategy of projection and hard clamping (Equation 1) may be thought of as an analogue of *relative colorimetric* intent, while *absolute colorimetric* intent for svBRDFs is not in general achievable, since the dimensionality of BRDF space is greater than the number of printer BRDFs available, hence most BRDFs are not reproducible exactly.


## 4.2 Halftoning Algorithms for BRDFs

We utilize a halftoning algorithm in our pipeline to convert the continuous-value printer weight maps into binary maps. Recall that we assume that the term “printer BRDF” may in fact refer to a layered combination of actual inks, or to the substrate itself. Therefore, the goal of the halftoning algorithm is to pick exactly one printer BRDF to be placed at each pixel.

In our experiments, we found that many of the more exotic inks, such as metal foils, impose additional constraints on the printing process. Because these inks were originally designed for “spot color” use, rather than high-frequency halftoning, they may only be successfully dithered at a relatively low spatial frequency (in practice, a block of  $3 \times 3$  or  $4 \times 4$  pixels at 600 dpi). Otherwise, we observe significant artifacts: either the foil does not stick to the paper, or too much foil is deposited. Unlike the similar case of “dot gain” in traditional printing, the effect is random and is not easily modeled. Therefore, we impose one additional constraint on the halftoning: each printer BRDF must be used in a contiguous spatial “cluster” of some minimum pixel count  $c$  (typically between 8 and 16 pixels).

Our halftoning implementation is inspired by the clustered Hilbert-curve error diffusion approach of Velho and Gomes [1991], but adapted for multi-channel halftoning (the original worked only on grayscale images) and strictly enforcing the minimum cluster size  $c$ . Our algorithm walks along a space-filling curve, maintaining a cumulative error of the halftoning approximation for each channel. That is, for each channel we store the accumulated intensity from the continuous maps *minus* the number of pixels of that channel that have been emitted. This number may be positive or negative, and we take its absolute value to determine how badly each channel is approximating the input maps: we strive to choose the output that will minimize the approximation error over *all* channels. We also maintain a “look-ahead” error, which stores what the approximation error will be in  $c$  pixels.

Our algorithm begins by looking ahead by  $c$  pixels, determining the worst-case approximation error (over all channels) resulting from every possible choice of output pixel. We begin outputting pixels for the channel that results in the lowest approximation error, and continue outputting that channel for at least  $c$  pixels. At that point, we examine the approximation errors again to decide whether to continue with the same channel (with no requirement that we output an exact multiple of  $c$  pixels: we must simply output  $\geq c$  contiguous pixels of the same channel) or to switch to a new channel. The algorithm is efficient, requiring only a single pass over the image, and typically takes only a second or two for a full page at 600 dpi.



**Figure 7:** Left: Halftoned weight maps for the four passes required to print the “oak leaf” svBRDF. Right: Closeup of clustered-dot error diffusion.

The very last step before printing is to combine the dithered maps for the “printer BRDFs” that we had been using, which may consist of multiple layers, to create maps for individual printing passes. The result of the halftoning pipeline for the oak leaf dataset is presented in Figure 7. This svBRDF was decomposed into four passes, utilizing the metallic silver, cyan, magenta, and yellow cartridges. Note that at this final stage the maps *can* overlap, if this is necessary to create the desired BRDF at some pixel.

**Discussion:** We chose a halftoning algorithm based on Hilbert-curve error diffusion chiefly for its ease of implementation, as well as its adaptability to multiple output channels and minimum cluster size. We believe that most other halftoning and dithering algorithms, whether ordered or based on error diffusion, could be adapted in similar ways. Moreover, future output technologies may eliminate the cluster-size constraints, allowing any multi-channel halftoning algorithm to be used. We expect the strengths and weaknesses of different algorithms, such as the tradeoffs between structured and random patterns, to carry over to the svBRDF setting.

## 5 Hardware Printing Setup

We require a printer capable of producing a wide range of materials in order to demonstrate the whole reflectance printing process. Commercial offset printers are capable of producing a wide range of different materials including metallic inks, glossy overcoats, etc. However, experimenting with this printing technology is expensive and time-consuming. On the other hand, the range of reflectances available with ink-jet technology is steadily growing (e.g., metallic inks are just being introduced), but it still does not span a range that is sufficiently interesting. Therefore, we have used an older printer model, the ALPS MD-5500 (also sold as the Okidata DP-5000) for validation of our pipeline. The main advantage of this printer is that it can print a wide range of different materials (diffuse, metallics, overcoats, foils, etc.), using a thermal-resin technology that places dots of pigment-impregnated resin onto the paper. The printer is relatively inexpensive (about 800 USD), and we have developed drivers based on the open-source `ppmtomd` package that allow flexible printing of multiple passes.

### 5.1 Inks

The ALPS MD-5500 uses a variety of different printer cartridges. Although at most 7 ink cartridges can be loaded at one time, cartridges can be swapped between passes. Therefore, it is possible to print a document that uses all available cartridges. In particular, we use the following 12 inks available from Okidata: cyan, magenta, yellow, black, metallic cyan, metallic magenta, metallic gold, metallic silver, gold foil, silver foil, finish, and primer. In



To perform white balancing and absolute calibration, a Gretag-Macbeth three-step gray card is placed in the scene, visible in every image. By using the same high-dynamic-range image reconstruction process we compute the average red, green and blue pixel values for the gray card and use them as a reference. To determine the absolute reflectance values we make use of the card’s middle (gray) part, which is designed to reflect 18% of incoming light.

## 6 Results

We have evaluated our processing pipeline on a number of different examples. In particular, we have used two svBRDF data sets from Lawrence et al. [2006]: “Dove” and “Season’s Greetings.” Furthermore, we have captured two additional svBRDF data sets: an “Oak Leaf” and a “Rusty Flange.” All of these data sets contain BRDFs that are not simply diffuse reflection.

The “Dove” data set is decomposed into 3 basis materials. It contains both metallic and diffuse BRDFs. As seen in Figure 6, the original svBRDF can be represented well with the linear combination of basis inks. Furthermore, the printed output comes quite close to the original material. The “Season’s Greetings” data set is decomposed into 4 basis materials. This data set is more challenging, since the materials span a large range of specular values. As can be seen in Figure 1, these linear combinations of basis inks come close to the original, but some inaccuracy relative to the original data is introduced. Nevertheless, despite the errors the printed output closely resembles the original sample. Both “Oak Leaf” and “Rusty Flange” are represented using 15 basis materials. When printed, both of these svBRDFs create a strong illusion of viewing an actual leaf or a rusty material.

Photographs of the Oak Leaf and Dove are presented in Figure 10, while Figure 1 shows a photo of the Season’s Greetings dataset. All are photographed under point lighting, though the curvature of the paper causes localized highlights to appear. For the Season’s Greetings and Rusty Flange datasets, we have obtained calibrated photographs under carefully controlled lighting conditions, enabling a comparison to the predicted renderings (Figure 11).

**Discussion:** The comparisons show the accuracy we are able to achieve, relative to our predictions, demonstrating that even though we might not always be able to match the original data because of the limited printer gamut, we can offer an accurate preview of the final printed document.


Due to the effective output resolution limit (only 150-200 dpi) of the ALPS printer, the printouts have to be viewed from at least half a meter distance. When viewed from closer distance, the half-toning patterns become visible and the material perception breaks down. We believe that using commercial offset printing technologies could alleviate this problem.

Another clear distinction from the original measured samples is the inability of our system to reproduce both normal and surface tangent variation, present in the original “Dove,” “Oak Leaf,” and “Rusty Flange” data sets.

## 7 Conclusions and Future Work

We have presented a framework for representing and printing documents as spatially-varying BRDFs, approximating the target svBRDFs with linear combinations of printer BRDFs. We have proposed a gamut-mapping algorithm for BRDFs that avoids clamping artifacts and preserves relative differences between reflectances in a printed document. Furthermore, our system allows us to preview documents with arbitrary lighting and view before they are printed. This feature can be especially useful for offset printing, which has higher set-up costs and requires larger runs.


In summary, we believe that our framework is an effective pipeline for printing documents with desired reflectance proper-



**Figure 10:** Results of printing svBRDFs using our pipeline, with visualizations of the target svBRDFs (under multiple light directions) for comparison. These are photographs under point-lighting, though a vertical highlight is effectively produced because of the curl of the paper. Note the significant material variation present in these examples: there are both diffuse areas and areas with significant gloss or specularity, such as the silver on the dove and highlighted regions on the oak leaf.

ties. This framework will enable creative professionals to control the complete reflectance properties of printed digital documents, smoothly extending upon the capabilities of color-only document-production pipelines.

Our current system has some limitations that suggest interesting directions for future work. First, we would like to take into account the scattering properties of the substrate. This will require representing documents using Bidirectional Surface Scattering Reflection Distribution Functions (BSSRDFs) rather than svBRDFs. This process would require measuring the scattering properties of the substrate, as well as more complex gamut-mapping and halftoning algorithms. Second, current printers can only print isotropic BRDFs. While it is natural to extend our software pipeline to anisotropic BRDFs, physically producing such materials will require using more sophisticated output devices that can embed mi-




**Figure 11:** Comparison of two target svBRDFs to physical printouts. The images of the printouts were captured at a perpendicular viewing direction and illuminated by a point light source at six different elevation angles. We estimated the precise camera and light locations along with their spectral output and sensitivity and the angular fall-off of the light source to produce renderings of the target svBRDF under nearly identical conditions. These comparisons demonstrate the high level of agreement we achieve between our printouts and the target reflectance.

crogeometry into printed surfaces. Third, we have only demonstrated printing on planar surfaces. It would be desirable to print spatially-varying BRDFs directly onto 3D objects (for example, objects produced by 3D printers) in order to design and control their reflectance properties. Finally, one could also improve our gamut-mapping algorithm, which globally remaps document reflectances. We believe that it is possible to design *local* gamut-mapping algorithms that locally preserve relative differences between document reflectances in a similar fashion to local tone-mapping operators for high-dynamic range images.

A further direction for future work is developing devices that allow *dynamically* changing between different BRDFs, leading to a true “BRDF display.” We believe that technologies such as e-ink are ideally suited for this application, since they would allow switching between different basis materials (with different BRDFs) at each pixel. It is also possible to use spatial multiplexing to accomplish the same goal, by using our current pipeline to print a checkered grid of different materials, then using an LCD panel to expose or conceal pixels of the grid. Figure 12 shows a prototype of such a device, using an off-the-shelf high-contrast LCD extracted from a black-and-white display. This prototype has some limitations, including a limited spatial resolution and the significant amount of light blocked by the LCD, yet suggests that it should be practical to create surfaces with dynamic spatially-varying materials, which would expand the possibilities of material design even more broadly.

## Acknowledgments

This work benefitted greatly from the comments of the anonymous referees, as well as conversations with colleagues at Adobe and elsewhere. The authors were partially supported by the US National Science Foundation, grants CCF-0347427, CCF-0702580, CNS-070820, CCF-0746117 and CCF-0747220. Thanks to J.C. Bradfield for making his ppmtomd package available, and to Paul Debevec for the environment maps used in natural-illumination renderings.



**Figure 12:** “BRDF display” based on a printed grid of materials, modulated by an LCD panel (left). By changing which portions of the underlying grid are exposed, we can display patterns in any combination of the basis BRDFs (visualized from two different viewpoints, center and right).



## References

- ASHIKHMIN, M., PREMOŽE, S., AND SHIRLEY, P. 2000. A Microfacet-Based BRDF Generator. In *Proc. ACM SIGGRAPH*, 65–74.
- COOK, R. L., AND TORRANCE, K. E. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graphics*, Vol. 1, No. 1, 7–24.
- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and Texture of Real-World Surfaces. *ACM Trans. Graphics*, Vol. 18, No. 1, 1–34.
- FLEMING, W., DROR, R. O., AND ADELSON, E. H. 2001. How Do Humans Determine Reflectance Properties under Unknown Illumination? In *Proc. CVPR Workshop on Identifying Objects Across Variations in Lighting: Psychophysics and Computation*, 347–368.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear Light Source Reflectometry. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, Vol. 22, No. 3, 749–758.
- GOLDMAN, D. B., CURLESS, B., HERTZMANN, A., AND SEITZ, S. M. 2005. Shape and Spatially-Varying BRDFs from Photometric Stereo. In *Proc. ICCV*, 341–348.
- HAN, C., SUN, B., RAMAMOORTHI, R., AND GRINSPUN, E. 2007. Frequency Domain Normal Map Filtering. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, Vol. 26, No. 3, 28:1–28:11.
- HERSCH, R. D., COLLAUD, F., AND EMMEL, P. 2003. Reproducing Color Images with Embedded Metallic Patterns. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, Vol. 22, No. 3, 427–436.
- HERSCH, R. D., DONZÉ, P., AND CHOSSON, S. 2007. Color Images Visible under UV Light. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, Vol. 26, No. 3, 75:1–75:9.
- KANG, H. 1999. *Digital Color Halftoning*. SPIE Press Monograph Vol. PM68.
- KAUTZ, J., AND MCCOOL, M. D. 1999. Interactive Rendering with Arbitrary BRDFs using Separable Approximations. In *Proc. Eurographics Workshop on Rendering*, 247–260.
- LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-Linear Approximation of Reflectance Functions. In *Proc. ACM SIGGRAPH*, 117–126.
- LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHI, R., AND RUSINKIEWICZ, S. 2006. Inverse Shade Trees for Non-Parametric Material Representation and Editing. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, Vol. 25, No. 3, 735–745.
- LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-Based Reconstruction of Spatial Appearance and Geometric Detail. *ACM Trans. Graphics*, Vol. 22, No. 2, 234–257.
- MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-Based BRDF Measurement Including Human Skin. In *Proc. Eurographics Workshop on Rendering*, 131–144.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A Data-Driven Reflectance Model. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, Vol. 22, No. 3, 759–769.
- MCALLISTER, D. K. 2002. *A Generalized Surface Appearance Representation for Computer Graphics*. Ph. D. thesis, University of North Carolina at Chapel Hill.
- MCCOOL, M. D., ANG, J., AND AHMAD, A. 2001. Homomorphic Factorization of BRDFs for High-Performance Rendering. In *Proc. ACM SIGGRAPH*, 171–178.
- MITSUNAGA, T., AND NAYAR, S. K. 1999. Radiometric Self Calibration. In *Proc. CVPR*, 374–380.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental Analysis of BRDF Models. In *Proc. Eurographics Symposium on Rendering*, 117–126.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2006. Image-Driven Navigation of Analytical BRDF Models. In *Proc. Eurographics Symposium on Rendering*, 399–408.
- OSTROMOUKHOV, V. 2001. A Simple and Efficient Error-Diffusion Algorithm. In *Proc. ACM SIGGRAPH*, 567–572.
- PANG, W.-M., QU, Y., WONG, T.-T., COHEN-OR, D., AND HENG, P.-A. 2008. Structure-Aware Halftoning. *ACM Trans. Graphics (Proc. ACM SIGGRAPH)*, Vol. 27, No. 3, 89:1–89:8.
- PELLACINI, F., FERWERDA, J. A., AND GREENBERG, D. P. 2000. Toward a Psychophysically-Based Light Reflection Model for Image Synthesis. In *Proc. ACM SIGGRAPH*, 55–64.
- PHONG, B. T. 1975. Illumination for computer generated pictures. *Comm. ACM*, Vol. 18, No. 6, 311–317.
- RUSINKIEWICZ, S. 1998. A New Change of Variables for Efficient BRDF Representation. In *Proc. Eurographics Workshop on Rendering*, 11–22.
- STOLLNITZ, E. J., OSTROMOUKHOV, V., AND SALESIN, D. H. 1998. Reproducing Color Images Using Custom Inks. In *Proc. ACM SIGGRAPH*, 267–274.
- TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for Off-Specular Reflection from Roughened Surfaces. *JOSA*, Vol. 57, No. 9, 1104–1114.
- ULICHNEY, R. 1987. *Digital Halftoning*. MIT Press, Cambridge, MA.
- VELHO, L., AND DE MIRANDA GOMES, J. 1991. Digital Halftoning with Space Filling Curves. In *Proc. ACM SIGGRAPH*, 81–90.
- WARD, G. J. 1992. Measuring and Modeling Anisotropic Reflection. In *Proc. ACM SIGGRAPH*, 265–272.

## Appendix BRDF Distance Metric

Given two BRDFs represented as the half-angle curves  $f_1(\theta_h)$  and  $f_2(\theta_h)$ , we use the following distance metric between them:

$$d(f_1(\theta_h), f_2(\theta_h)) \equiv \left( \int_{\theta_h=0}^{\pi/2} \|g_1(\theta_h) - g_2(\theta_h)\|^2 d\theta_h \right)^{1/2}, \quad (2)$$

$$\text{where } g(\theta) = (f(\theta) \cos \theta)^{1/3} \sqrt{\sin \theta \cos \theta}. \quad (3)$$

This is derived by considering spheres with the given BRDFs, illuminated by a directional light and observed from the light direction. For a given angle of incidence  $\theta$ , the irradiance is just  $f(\theta) \cos \theta$ , where we use the fact that, for co-incident lighting and viewer directions,  $\theta_i = \theta_r = \theta_h$ . We take the cube root of this to incorporate a perceptually-based nonlinearity, inspired by the nonlinearity present in the CIELAB color space. Now, we wish to take the  $L^2$  distance between the (nonlinearly-adjusted) sphere images, which requires us to account for the relative projected areas for each  $\theta$ . This is simply  $\sin \theta \cos \theta$ . However, this factor would ordinarily go outside the squared norm in equation 2, so in order to fold it into  $g$  (hence allowing it to be precomputed) we must take its square root in equation 3.