

Printing Spatially-Varying Reflectance for Reproducing HDR Images

Yue Dong*

Xin Tong*

Fabio Pellacini^{†‡}

Baining Guo*

*Microsoft Research Asia

[†]Dartmouth College

[‡]Sapienza University of Rome

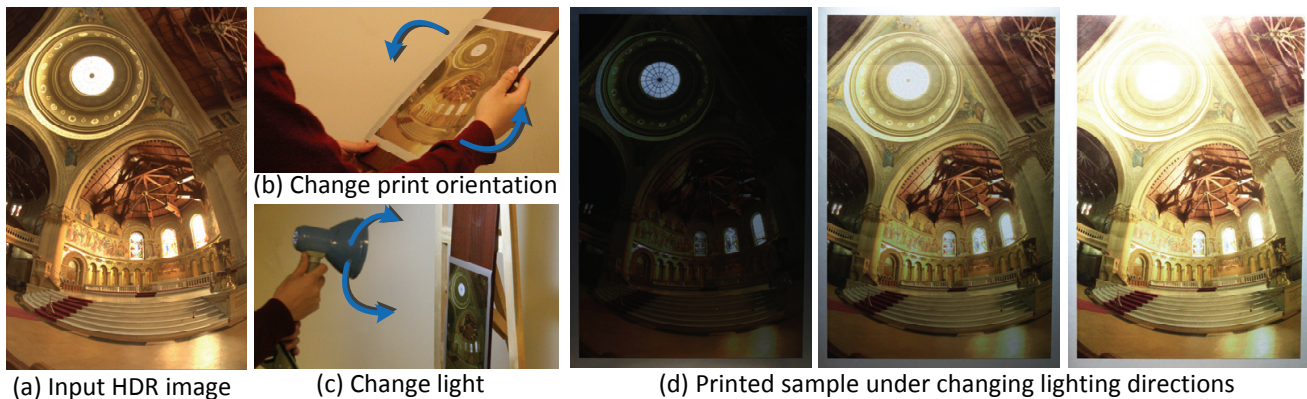


Figure 1: Given an HDR input image (a), we print the image with glossy inks in such a way that when either the print orientation (b) or the lighting direction (c) are changed, different luminance levels of the original image are displayed. The appearance of the print under changing lighting is shown in (d), while the appearance of the print under orientation change is shown in the accompanying video. Another result for viewing scenario (b) can be found in Figure 5. All results shown in this paper are photographs of actual prints. HDR image courtesy of Paul Debevec.

Abstract

We present a solution for viewing high dynamic range (HDR) images with spatially-varying distributions of glossy materials printed on reflective media. Our method exploits appearance variations of the glossy materials in the angular domain to display the input HDR image at different exposures. As viewers change the print orientation or lighting directions, the print gradually varies its appearance to display the image content from the darkest to the brightest levels. Our solution is based on a commercially available printing system and is fully automatic. Given the input HDR image and the BRDFs of a set of available inks, our method computes the optimal exposures of the HDR image for all viewing conditions and the optimal ink combinations for all pixels by minimizing the difference of their appearances under all viewing conditions. We demonstrate the effectiveness of our method with print samples generated from different inputs and visualized under different viewing and lighting conditions.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture;

Keywords: appearance fabrication, high dynamic range images

Links: [DL](#) [PDF](#)

ACM Reference Format

Dong, Y., Tong, X., Pellacini, F., Guo, B. 2012. Printing Spatially-Varying Reflectance for Reproducing HDR Images. *ACM Trans. Graph.* 31 4, Article 40 (July 2012), 7 pages. DOI = 10.1145/2185520.2185536 <http://doi.acm.org/10.1145/2185520.2185536>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2012 ACM 0730-0301/2012/08-ART40 \$15.00 DOI 10.1145/2185520.2185536 <http://doi.acm.org/10.1145/2185520.2185536>

1 Introduction

Capturing high dynamic range images (HDR) has become ubiquitous thanks to recent advances in imaging devices and reconstruction algorithms. Although techniques have been developed for showing HDR content on active displays [Seetzen et al. 2004; Wetzstein et al. 2011], reproducing HDR images in prints remains challenging due to the low dynamic range (LDR) of physical reflective media. The basic principle that is followed today is to attempt to map between the dynamic range of the input image and output print. The most common method is to compress the dynamic range of the HDR image using one of the many tone mapping methods available (see [Reinhard et al. 2010] for a recent review). Although these methods can preserve the overall visual appearance of the input, the original contrast of the HDR image is lost after printing. More recently, methods have been proposed to extend the dynamic range of print media, for example by using a projector aligned with the image [Bimber and Iwai 2008] or a stack of images illuminated by a light source [Ward 2002; Wetzstein et al. 2011]. These methods are limited in their use since specialized setups are needed to view the print.

In this paper, we display HDR image contents with spatially-varying distributions of glossy materials. We modulate the reflectivity of the print in such a way that a viewer can see different exposures of the HDR image by changing the print orientation or lighting direction. Figure 1 shows one of our printed samples viewed under a moving light. Our method differs substantially from prior work in that we do not attempt to map the dynamic range of the input image to the output print for a single view. We instead conceptually map different luminance levels of the HDR image into a continuous sequence of LDR images that are reproduced by the appearance of a glossy print when viewed under gradually rotating print orientations or lighting directions. Compared to prior work, the main advantage of our method is that we maintain the tonal details of the HDR image at all exposures, while requiring viewing conditions similar to standard prints. Furthermore, viewers can see the whole range of exposures generated by an HDR image by




Figure 2: Using the appearance of glossy inks to display high dynamic range content. (a) The appearance of the input HDR image under two different exposures. (b) The intensity variations of four image pixels under different exposures. The positions of the pixels are marked in (a). (c) The appearance variations of different glossy ink combinations under changing lighting directions. (d) A swatch of all valid ink combinations. Without optimization, the plots shown in (b) cannot match the plots in (c).

quickly changing the print orientation or lighting direction.

We print our images using commercially available printers that allow us to customize the print reflectance with ink combinations printed over metallic substrates. Unlike [Matusik et al. 2009], which is designed for approximating target BRDFs with the BRDFs of ink combinations, our goal is to use the BRDFs of inks to reproduce the appearance of an input HDR image under each viewing condition. As illustrated in Figure 2, because the intensity variations of image pixels under an exposure setting may not follow the appearance variation of BRDFs and the number of valid ink combinations is large, it is non-trivial to choose exposure settings and ink combinations so that the appearance of the HDR image and the ink combination match. Our main technical challenge is therefore to simultaneously determine the exposure of the HDR image for each viewing condition (i.e. the varying viewing and lighting directions of the print as determined by the viewer) and the ink distribution of all pixels that best reproduces the HDR image appearance with this exposure. We pose this as an optimization problem and solve for ink combinations and exposures iteratively from an initial guess.

We start by determining the range of exposures that best captures the tonal range of the HDR image. We evenly assign exposures in this range to the continuous viewing conditions. This initial assignment might not be optimal since the chosen exposures may not be easily reproducible by the available inks. We thus iteratively refine both the ink combination and exposures. In each iteration, we first determine the ink combinations such that the printed appearance best matches the exposures of the HDR image when viewed at the corresponding viewing conditions. Then we refine the exposures of the HDR image to best match the appearance of the current ink combinations displayed at the corresponding viewing conditions. We iterate these two steps until there is no more improvement.

The prints generated by our method can be viewed in a wide range of conditions, just like standard prints, without the help of dedicated setups. At the same time, we are able to well preserve the visual contents of the input image by allowing the viewer to see different exposures simply by changing the print orientation or light direction. The remainder of this paper will describe our printing process and show several results of actual printed materials.

2 Related Work

A comprehensive survey of techniques to capture, display, and render HDR images can be found in [Reinhard et al. 2010]. In this section, we review the work most closely related to our own.

HDR display and printing Many methods have been proposed to extend the dynamic range of displays by overlaying two or more translucent light modulators to multiply their contrast. Ward [2002] decomposes an HDR image into two LDR images that are printed

as transparencies and stacked together. The original HDR image is viewed by aligning the two layers in front of a bright backlight. Seetzen et al. [2004] combine an LCD with either a DLP projector or an LED panel to make an HDR display. Wetzstein et al. [2011] extends these methods using multiple stacked transparent slides for displaying both 4D light fields and HDR images. Bimber and Iwai [2008] extend the dynamic range of reflective media with a calibrated projector. Instead of increasing the dynamic range of the output display, we visualize different exposures of the original HDR at different views, requiring no specialized viewing hardware.

Early artistic work used glossy material to enhance the brightness of prints [Reinhard et al. 2010]. However, these prints can only be viewed under specific lighting and viewing directions. In contrast, our method exploits the angular variation in the appearance of glossy surfaces to display different exposures of HDR images. Furthermore, our method works well for a wide range of viewing conditions.

Appearance fabrication A set of methods has been developed for fabricating objects with custom surface reflectance [Weyrich et al. 2009; Matusik et al. 2009; Hullin et al. 2011], subsurface scattering [Dong et al. 2010a; Hašan et al. 2010], and deformation behaviors [Bickel et al. 2010]. Weyrich et al. [2009] reproduce a custom surface reflectance using a microfacet pattern fabricated on a reflective physical surface. Matusik et al. [2009] use a set of glossy inks for printing spatially-varying BRDFs. Hullin et al. [2011] propose a solution for displaying dynamic BRDFs with liquid surfaces. Our method follows Matusik et al.’s work [2009] in its use of glossy inks for printing. But since the goals of the two methods differ substantially, we propose a new optimization method for determining both ink combinations and exposure settings.

A method close to ours is the one of Alexa and Matusik [2010] that constructs reliefs that display different images when viewed under different directional illuminations. Although their method can reproduce a discrete number of images, it is difficult to extend to continuous image sequences required in our application. Furthermore, the spatial resolution of the result is limited. Most recently, Holroyd et al. [2011] developed a method for converting a digital 3D model into a fabricated multilayer model. Although this approach displays the object shapes for a wide range of viewing directions, the surface texture is printed as diffuse colors and remains constant when changing views.

3 Printing Spatially Varying BRDFs for HDR Images

Given an input HDR image, we print combinations of glossy inks so that the viewer can see different exposures of the original image when rotating either the print or the lighting. In this section, we




Figure 3: Ink swatches and the BRDF gamut measured from ink swatches. Note that the ink swatches exhibit different appearance variations under different lighting directions due to the BRDFs of different ink combinations. Boundary regions without ink illustrate the appearance of the metallic substrate used for the print. The gamut of the red channel of the BRDFs is illustrated in (b).

first describe our printing solution and then focus on how we determine the exposures and ink combinations necessary to reproduce the desired effect.

3.1 Printing Hardware

Our method is based on printing hardware that supports the Process Metallic Color System [Color-Logic 2011]. While conventional color printers are based on CMYK inks, these printers use an additional silver ink when using a white paper substrate, or a white ink when using a metallic substrate. These printers are capable of controlling the amount of glossiness with the 8-bit “color” depth of the additional ink. The results in this paper are printed with an HP Indigo 5500 on a flat metallic substrate with CMYK and white inks. The underlying metallic substrate is homogeneous and has a fixed isotropic BRDF. Each ink combination, which is represented by color and white ink tuples $c = (r, g, b, w)$, determines a resulting printed BRDF $\rho_c(x)$. We express our color in RGB since the printer’s driver automatically converts them to CMYK. Figure 3(a) illustrates the appearances of ink swatches printed on a metallic substrate under two lighting directions.

The BRDF gamut is composed by the set of BRDFs ρ_c that correspond to all valid ink combinations $c \in \mathcal{C}$. We model the gamut by assuming that the gamut’s BRDFs lie in a low dimensional manifold in the BRDF space [Matusik et al. 2003; Dong et al. 2010b]. We thus represent each BRDF as a locally-linear combination of BRDFs with known ink combinations. We choose as a basis the BRDFs of a sparse set of ink combinations \mathcal{C}_K , which are measured using linear light source reflectometry [Gardner et al. 2003]. The BRDF ρ_c corresponding to the ink combination $c = \sum_{k=1}^M \omega_k c_k$, can be computed as $\rho_c = \sum_{k=1}^M \omega_k \rho_{c_k}$, where ρ_{c_k} are the basis BRDFs for the ink combinations $c_k \in \mathcal{C}_k$ in the local neighborhood of size $M = 5$.

In practice, we uniformly sampled 16 levels for each of the input RGB color channels and the additional white ink channel, resulting in sixty-five thousand different ink combinations. We tiled all ink combinations in a 5120×5120 image with a 20×20 patch for each ink combination and printed the image on an A4 metallic substrate. We then measured the BRDFs of all surface points via linear light source reflectometry [Gardner et al. 2003] and computed the BRDF of each sampled ink combination by averaging the BRDFs in the corresponding tile. With this density the locally-linear model can predict well the printed BRDF for a given ink combination. In our implementation, we represent the BRDF of a sampled ink combination with the isotropic Ward model. Figure 3 (b) illustrates the BRDF gamut of the red channel of all ink combinations, where the roughness of the BRDFs increases as the intensity of white ink increases.




Figure 4: The optimization pipeline. Given an input HDR image $I_h(x)$, our method first initializes the exposure settings for each frame t and then optimizes the ink combinations and exposure settings iteratively. We finally print out a physical sample based on the computed ink combinations for all the pixels.

3.2 Viewing Conditions

We view the print in an environment with ambient illumination and one bright directional light source. Since the distance between the viewpoint and the print is large considering the size of the print and the sharpness of the specular lobe, we assume that the viewing direction for all points on the print is the same. As shown in Fig. 1, our system supports two kinds of viewing scenarios. In the first, the print is fixed and the user rotates the light source around the print. In the second, the user holds the print and rotates it under a fixed view and lighting. Due to the reciprocity of the BRDF, similar results can be achieved with either moving the view or the light. Figure 9 displays appearance variations of one physical print under two viewing scenarios.

For both scenarios, we parameterize the viewing sequence by $0 \leq t \leq T$ and model the viewing conditions $V(t) = \{v(t), E(t)\}$ as the set of viewing directions $v(t)$ and environmental lighting $E(t)$ for the sequence t . Both $v(t)$ and $E(t)$ are defined with respect to the local frame of the print. The environmental lighting is either measured from real viewing environments with a light probe or we use a default setting that consists of a white ambient light and a white directional light with a radiance ratio of 1 : 100. In practice, we assume that the half angle (i.e. the direction halfway between the viewing and lighting directions) varies from 0 degrees (specular peak direction) to 45 degrees. We uniformly sample the viewing sequence with $T = 20$ samples for optimization.

3.3 Exposure and Ink Combination Optimization

Given an input HDR image $I_h(x)$, we compute the ink combination $c(x)$ of each pixels x and the exposure setting $s(t)$ of all viewing conditions $V(t)$ by minimizing the difference between the appearance of the print and appearance of the HDR image at the chosen exposures for the corresponding viewing conditions:

$$\arg \min_{c(x), s(t)} \sum_x \sum_t \|I_r(\rho_{c(x)}, V(t)) - I_s(I_h(x), s(t))\|^2, \quad (1)$$

where $\rho_{c(x)}$ is the BRDF of the ink combination at x . $I_r(\rho_{c(x)}, V(t))$ is the appearance of the print under the viewing condition $V(t)$, which is rendered with the given environmental lighting $E(t)$ and viewing direction $v(t)$ as

$$I_r(\rho_{c(x)}, V(t)) = \int_{\Omega} \rho_{c(x)}(v(t), l_i) (N \cdot l_i) E(t, l_i) d\omega_{l_i}, \quad (2)$$

where l_i is the lighting direction and $E(t, l_i)$ is the incident radiance of environmental lighting along direction l_i . We normalize the

rendering result to the range of 0 to 1, picking the maximum radiance that can be generated by the BRDFs in the gamut under the given viewing conditions.

$I_s(I_h(x), s(t))$ is the appearance of the input image $I_h(x)$ at the exposure setting $s(t)$ for viewing condition $V(t)$, which is computed by

$$I_s(I_h(x), s(t)) = \min\{2^{s(t)} I_h(x) / \bar{I}_w, 1.0\}, \quad (3)$$

where \bar{I}_w is the "key" value of HDR image $I_h(x)$ and computed by $\bar{I}_w = \exp(\frac{1}{N} \sum_x \log(1.0e - 5 + I_h(x)))$ [Reinhard and Devlin 2005]. N is the total number of pixels in the image. We clamp all computed image values to the range of 0 to 1. Note that with the exponential exposure setting used in Equation 3, the logarithmic response properties of human perception have been taken into account in the optimization.

As illustrated in Figure 4, we solve this optimization problem iteratively from an initial guess for $s(t)$. In each step, we first fix the exposures $s(t)$ to solve for the ink combinations $c(x)$. We then fix the ink combinations $c(x)$ to update the exposures $s(t)$. We repeat these two steps until the image differences described in Equation 1 are not reduced anymore.

Initialization To initialize the exposure settings of the HDR image for each t , we chose a desired range of exposures (s_{min}, s_{max}) that well captures the tonal range of the HDR $I_h(x)$. For our prints, we set the minimal exposure s_{min} in such a way that the non-zero darkest pixels I_h^{min} are mapped to 0.1, while the maximal exposure s_{max} is set to map the brightest image pixels I_h^{max} to 0.9. According to Equation 3, we have $s_{min} = \log_2(0.1 \bar{I}_w / I_h^{min})$ and $s_{max} = \log_2(0.9 \bar{I}_w / I_h^{max})$. After that, we evenly assign the exposures in the range to all viewing conditions to get the initial exposure $s(t) = s_{min} + \frac{t}{T}(s_{max} - s_{min})$ for each viewing condition. Figure 5(a) shows the appearance of an input image at the initial exposure setting for one viewing condition.

Ink Combination Optimization Given the exposure settings $s(t)$ for the viewing conditions $V(t)$, we first compute the corresponding input image appearances $I_s(x, t) = I_s(I_h(x), s(t))$ with Equation 3 and then use these images to optimize the ink combination $c(x)$ at each pixel by computing

$$\arg \min_{c(x)} \sum_t \|I_r(\rho_{c(x)}, V(t)) - I_s(x, t)\|^2. \quad (4)$$

Figure 5(c) and (f) illustrate this optimization, where the dotted lines plot the intensity values of two pixels $I_s(x, t)$ for all viewing conditions, while the solid lines are appearance variations of two pixels $I_r(\rho_{c(x)}, V(t))$ rendered by their BRDFs under all viewing conditions.

Since the BRDF $\rho_{c(x)}$ in the print gamut is represented as the locally-linear combination of basis BRDFs $\rho_{c(x)} = \sum_{k=1}^M \omega_k(x) \rho_{c_k(x)}$, we have $I_r(\rho_{c(x)}, V(t)) = \sum_{k=1}^M \omega_k(x) I_r(\rho_{c_k(x)}, V(t))$. Therefore, we solve for the ink combination by searching for a locally-linear combination of basis BRDFs that can best reproduce the pixel's appearance under all corresponding viewing conditions:

$$\arg \min_{\omega_k(x), \rho_{c_k(x)}} \sum_t \left\| \sum_{k=1}^M \omega_k(x) I_r(\rho_{c_k(x)}, V(t)) - I_s(x, t) \right\|^2. \quad (5)$$

To this end, we first search for M basis BRDFs whose appearances under the given viewing conditions best match the input image exposures $I_s(x, t)$. We accelerate this process by precomputing the

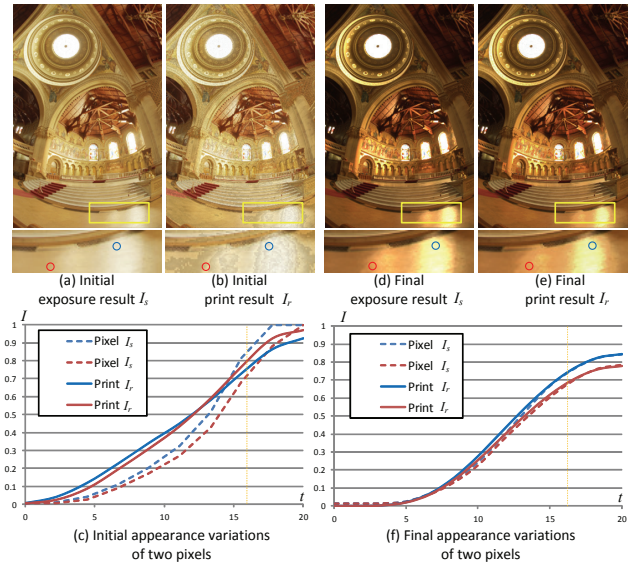


Figure 5: Iterative optimization results at a specific viewing condition. The initial exposures of input image (a) cannot be well reproduced by appearance of ink combinations (b). As shown in (c), the intensity values of two pixels at different exposures (dotted line) is different from the appearance variations generated by ink combinations (solid line). Note that in the zoomed region, the relative contrast of the two pixels (marked as circles) in the initial exposure result is even inverted in the print result and leads to an artifact. After the iterative optimization, the exposure of the input (d) and the appearance of the print (e) become consistent. The appearances of the two pixels match (f) and the artifact in the zoomed region is removed.

appearance of all basis BRDFs under the given viewing conditions and then searching for the results via Approximate Nearest Neighbor (ANN) search [Mount and Arya 1997]. Given the local basis BRDFs $\rho_{c_k(x)}$ and their appearances $I_r(\rho_{c_k(x)}, V(t))$, we then derive the weights of the local basis BRDFs $\omega_k(x)$ by solving the linear system determined by Equation 5. Finally, the ink combination at x is computed as the weighted sum of the ink combinations associated with the local basis BRDFs $c(x) = \sum_{k=1}^M \omega_k(x) c_k(x)$. In the case that the resulting ink combination is outside of the of print gamut, we simply approximate the result with the closest ink combination in the gamut. Figure 5(b) and (e) illustrate the images $I_r(\rho_{c(x)}, V(t))$ of the ink combinations optimized from the initial exposure settings (Figure 5(a)) and the image of final ink combinations after iteration, both of which are rendered under the same viewing conditions.

Our method does not include explicit spatial coherency constraints in ink combination optimization. Instead, we assume that for pixels with similar intensities, their similar appearances under the same exposure settings would result in similar ink combinations. For two pixels with small intensity variations, the differences of the two resulting BRDF basis sets (in BRDF space) are small due to the dense samples of the BRDF bases. By fitting the appearance of two pixels by local linear combinations of BRDF bases, our method guarantees that the differences of the resulting BRDFs of two pixels are very small. In principle, different ink combinations may be derived for pixels with the same intensity and thus generate artifacts in the final result. Nevertheless, we have not encountered such issues in practice.

Exposure Optimization Given the ink combinations $c(x)$ and their appearance $I_r(x, t) = I_r(\rho_{c(x)}, V(t))$, we optimize the exposure settings $s(t)$ of the input image for each viewing condition t by computing

$$\arg \min_{s(t)} \sum_x \|I_r(x, t) - I_s(I_h(x), s(t))\|^2. \quad (6)$$

By setting $s'(t) = 2^{s(t)}$ and $I_s(I_h(x), s(t)) = s'(t)I_h(x)/\bar{I}_w$, we convert the optimization to a least squares fitting problem:

$$\arg \min_{s(t)} \sum_x \|I_r(x, t) - s'(t)I_h(x)/\bar{I}_w\|^2, \quad (7)$$

which can be directly solved by $s'(t) = \frac{1}{N} \sum_x \frac{I_r(x, t)\bar{I}_w}{I_h(x)}$. The exposure of the input image for each viewing condition is then determined by $s(t) = \log_2(s'(t))$. Figure 5(d) shows the image of the input rendered at the optimized exposure setting.

Iterative Optimization Since the optimization goal in Equation 1 is affected by both the ink combinations $c(x)$ and the exposure settings $s(t)$, the iterative optimization is required to find the optimal solution for both $c(x)$ and $s(t)$ while avoiding artifacts in the final print result. In the ink combination optimization step, the updated exposures constrains the spatial relationship of the ink combinations $c(x)$ to match the image content. While in the exposure optimization step, the updated monotonic appearance variation of BRDFs determined by the ink combination $c(x)$ at each pixel implicitly constrains the resulting exposures $s(t)$ for different viewing conditions. As a result, in the physical prints the exposure can vary monotonically under different viewing conditions while the relative contrast of different HDR regions is preserved well. Figure 5 illustrates the appearance of the print and the exposure of the input image under a specific viewing condition after the first iteration and for the final optimization result, showing the benefits of this iterative process.

4 Results

Implementation We ran our optimization on a PC with an Intel Xeon 2.83GHz CPU and 8GB RAM. For a typical HDR image at the resolution of 768×768 , the overall optimization time is less than 10 seconds. After optimization, the result is printed with an HP Indigo 5500 printer. Each image is printed on an A4 metallic substrate at the resolution of 700DPI. The average printing time is less than one minute for an A4 size print.

All the results shown in the paper and video are captured with a Canon EOS 5D Mark II with an EF 100mm F2.8 Lens. We calibrated its color using the standard color checker for all viewing conditions. The camera exposure (i.e. ISO, shutter speed, and aperture) is fixed during capture of each result so that the appearance variations of each print under different viewing conditions can be shown. For both viewing scenarios, we kept the lighting direction of the dominant light, the viewing direction, and the normal of the print roughly in the same plane and rotated the print orientation or light within the plane.

Results Figure 6(b) illustrates the intensity of white ink used for the print shown in Figure 1(a). In our solution, the brighter pixels are printed with a higher intensity of white ink, which produces BRDFs with wider specular lobes. On the contrary, the ink combinations with a low intensity of white ink are used for generating BRDFs with sharp specular lobes, which are used for printing low intensity pixels. Figure 6(c) illustrates the subset of BRDFs used for this input.

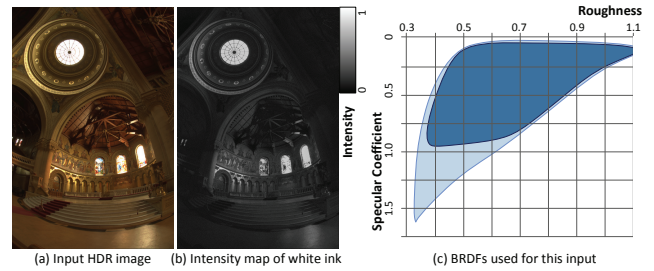


Figure 6: Ink combination used for printing. (a) Input HDR image and the resulting white ink intensity used for printing. (b) Subset of BRDFs (shown in dark blue) used for printing the input image.




Figure 7: Direct printing of HDR content in a single view. Although high dynamic range can be observed at certain view and lighting conditions (a) and (b). Slight changes in the light direction results in an unsatisfactory viewing experience (c).

Figure 8 shows the appearances of the same print viewed in two different lighting environments, in which we rotated the print to display the exposures of the HDR image. The lighting environment of the dark room is similar to our default setting while the lighting in the office environment is dominated by a long area light source on the ceiling. Despite the difference of the lighting environment, the resulting print shows the exposures of input images in both environments. Figure 9 shows another print viewed in two different viewing scenarios. Our method works well for both viewing scenarios with the same print.

Figure 10 compares the appearance of the print generated by our method and the LDR images that are generated from the same input image with different tone mapping operators [Fattal et al. 2002; Durand and Dorsey 2002; Reinhard et al. 2002] and used for conventional printing. All three LDR images are directly copied from the original papers. While the conventional printing solution that is based on the LDR images tries to map all local contrast into one image, our method illustrates image details at different exposures under different lighting directions.

Since the glossy substrate and rich ink combinations used in our solution provides higher contrast and more detailed intensity levels, it is possible to directly print the HDR image. Figure 7 shows the print result of an HDR image printed by this approach via the same printing hardware. Although the print could reproduce the high dynamic range effects of the input image at the specific mirror reflection direction, its appearance quickly changes and exhibits artifacts as the lighting or viewing direction slightly changes. As shown in Figure 7(b), the contrast between the sky and the building are inverted. On the contrary, our method well reproduces the appearance of the HDR input over the entire view range as shown in Figure 8.

Figures 11 and 12 show the appearance of two prints with rich tonal details. Under different lighting directions, the prints generated by our method well illustrate all these tonal details. Please see the accompanying video for continuous visual effects of all print results under different viewing scenarios.

Limitations Our method has several limitations. First, due to the limited BRDF gamut of the printing system, our method may print pixels with different luminance values using the same ink combination and thus fail to reproduce the tonal details of the input image at some exposures. Our method assumes that the environment includes only one directional light and that the contrast between the light source and the environment is high. In low dynamic range environments or environments that include multiple light sources, the visual effects of the print will degrade. Finally, although our method can display the exposures of images at different viewing conditions, the dynamic range of the print is still low and thus limits the tonal details that can be displayed for each view.

5 Conclusions

In this paper, we present a printing solution for reproducing different exposures of HDR images with spatially-varying surface reflectance. Our method exploits the appearance variation of the surface reflectance to display different exposures of the HDR image under varying viewing or lighting directions. Our method uses relatively inexpensive commercially available printers, providing a practical solution for displaying HDR images with low dynamic range reflective media. Compared to existing methods, we believe that our method provides an alternative solution for HDR image printing and enriches the choices of users in many applications.

A promising direction for future work is to extend the print gamut of our solution by using different printing systems that support different kinds of metallic inks, such as the one used in [Matusik et al. 2009]. It would also be interesting to investigate how to combine surface reflectance and normal variations to enhance the dynamic range of the print. Finally, how to leverage HDR and material perception for better predicting the visual effects of the final print under different lighting environments is another interesting research direction.

Acknowledgments

The authors would like to thank Stephen Lin for paper proofreading. The authors also thank the anonymous reviewers for their helpful suggestions and comments. Fabio Pellacini was supported by the NSF (CNS-070820, CCF-0746117), Intel and the Sloan Foundation.

References

- ALEXA, M., AND MATUSIK, W. 2010. Reliefs as images. *ACM Trans. Graph.* 29 (July), 60:1–60:7.
- BICKEL, B., BÄCHER, M., OTADUY, M. A., LEE, H. R., PFISTER, H., GROSS, M., AND MATUSIK, W. 2010. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph.* 29 (July), 63:1–63:10.
- BIMBER, O., AND IWAI, D. 2008. Superimposing dynamic range. *ACM Trans. Graph.* 27 (December), 150:1–150:8.
- COLOR-LOGIC, 2011. The color-logic process metallic color system.
- DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph.* 29 (July), 62:1–62:10.
- DONG, Y., WANG, J., TONG, X., SNYDER, J., LAN, Y., BEN-EZRA, M., AND GUO, B. 2010. Manifold bootstrapping for svbrdf capture. *ACM Trans. Graph.* 29 (July), 98:1–98:10.
- DURAND, F., AND DORSEY, J. 2002. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Trans. Graph.* 21 (July), 257–266.
- FATTAL, R., LISCHINSKI, D., AND WERMAN, M. 2002. Gradient domain high dynamic range compression. *ACM Trans. Graph.* 21 (July), 249–256.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. In *ACM SIGGRAPH 2003 Papers*, ACM, New York, NY, USA, SIGGRAPH '03, 749–758.
- HAŠAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph.* 29 (July), 61:1–61:10.
- HOLROYD, M., BARAN, I., LAWRENCE, J., AND MATUSIK, W. 2011. Computing and fabricating multilayer models. *ACM Trans. Graph.* 30 (Dec.), 187:1–187:8.
- HULLIN, M. B., LENSCH, H. P. A., RASKAR, R., SEIDEL, H.-P., AND IHRKE, I. 2011. Dynamic display of BRDFs. In *Computer Graphics Forum (Proc. EUROGRAPHICS)*, Blackwell, Llandudno, UK, O. Deussen and M. Chen, Eds., Eurographics, 475–483.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Trans. Graph.* 22 (July), 759–769.
- MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H. P. A., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing spatially-varying reflectance. *ACM Trans. Graph.* 28 (December), 128:1–128:9.
- MOUNT, D., AND ARYA, S. 1997. Ann: A library for approximate nearest neighbor searching. In *CGC 2nd Annual Fall Workshop on Computational Geometry*.
- REINHARD, E., AND DEVLIN, K. 2005. Dynamic range reduction inspired by photoreceptor physiology. *Visualization and Computer Graphics, IEEE Transactions on* 11, 1 (jan.-feb.), 13–24.
- REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. *ACM Trans. Graph.* 21 (July), 267–276.
- REINHARD, E., WARD, G., PATTANAIK, S., DEBEVEC, P., HEIDRICH, W., AND MYSZKOWSKI, K. 2010. *High Dynamic Range Imaging: Acquisition, Display, and Image-based Lighting*, 2nd ed. The Morgan Kaufmann series in Computer Graphics. Elsevier (Morgan Kaufmann), Burlington, MA.
- SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. *ACM Trans. Graph.* 23 (August), 760–768.
- WARD, G. 2002. A wide field, high dynamic range, stereographic viewer. In *PICS*, 30–34.
- WETZSTEIN, G., LANMAN, D., HEIDRICH, W., AND RASKAR, R. 2011. Layered 3d: tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Trans. Graph.* 30 (August), 95:1–95:12.
- WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (Aug.).




Figure 8: The appearances of a print viewed under different lighting environments. The ones on the left are captured in a dark room while the ones on the right are captured under office lighting. The environmental lightings of these two environments are shown in the middle. Both results clearly show the exposures of the input image under different print orientations. The main visible difference is due to the different lighting colors and the fixed camera white balance used in this comparison. HDR image courtesy of Paul Debevec.



Figure 9: The appearances of a print viewed under both viewing scenarios. HDR image courtesy of Industrial Light and Magic.



Figure 10: Comparisons between the traditional print solution and our method. The tone mapped LDR images used for conventional diffuse print (shown on the left) attempt to preserve local contrast in a single view. Our print shown on the right displays different exposures for different lighting directions. HDR image courtesy of Jack Tumblin.



Figure 11: The appearances of a print under different lighting directions for an image of an illuminated store at night. The print generated by our method well illustrates all tonal details at different exposures. HDR image courtesy of Greg Ward.



Figure 12: The appearances of a print of a Napa valley landscape image. The details in the sky and the ground regions are illustrated under different lighting directions. HDR image courtesy of Spheron.