EUROGRAPHICS 2012 / P. Cignoni, T. Ertl
(Guest Editors)

Volume 31 (2012), Number 2

Manufacturing Layered Attenuators for Multiple Prescribed Shadow Images

Ilya Baran[1] Philipp Keller[2] Derek Bradley[1] Stelian Coros[1] Wojciech Jarosz[1] Derek Nowrouzezahrai[1] Markus Gross[1,2]

[1]Disney Research Zürich [2]ETH Zürich

Figure 1: *A single manufactured multi-layer attenuator, consisting of three transparencies separated by acrylic, casts different shadows depending on the angle of the sun. The shadows match the target Van Gogh paintings, shown as insets.*

Abstract

We present a practical and inexpensive method for creating physical objects that cast different color shadow images when illuminated by prescribed lighting configurations. The input to our system is a number of lighting configurations and corresponding desired shadow images. Our approach computes attenuation masks, which are then printed on transparent materials and stacked to form a single multi-layer attenuator. When illuminated with the input lighting configurations, this multi-layer attenuator casts the prescribed color shadow images. Alternatively, our method can compute layers so that their permutations produce different prescribed shadow images under fixed lighting. Each multi-layer attenuator is quick and inexpensive to produce, can generate multiple full-color shadows, and can be designed to respond to different types of natural or synthetic lighting setups. We illustrate the effectiveness of our multi-layer attenuators in simulation and in reality, with the sun as a light source.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture I.4.0 [Image Processing and Computer Vision]: General—Image Displays

1. Introduction

Shadows are an integral part of our visual experience, providing depth cues and conveying information about the reflectance and lighting of the surrounding environment. Humans have been manipulating and using shadows throughout history, creating shadow puppets for entertainment, or building sundials to keep time. Transparent occluders, such as stained-glass windows, can cast colored shadows and can be used to produce desired images, as in a slide projector.

In this work, by putting computation in the loop, we combine the ideas of a slide and a sundial: we describe how to construct *layered attenuators*, whose shadows form different images depending on the lighting configuration (Figure 1). Layered attenuators are seemingly random patterns printed on transparencies and stacked together, which, when properly illuminated, cast different color images as shadows. Layered attenuators are easy to produce using commodity hardware and inexpensive: even our prototypes cost under $10 each. We demonstrate that they work in both natural

and synthetic lighting and do not require precise positioning. Surprisingly, layered attenuators can produce more distinct images than the number of layers.

The input to our algorithm is a list of target shadow images and corresponding light configurations (e.g., sun positions). Computing the patterns to be printed for the layered attenuator to obtain the desired shadows is an inverse light transport problem. Tomography approaches, which have been applied to related problems, break down for highly dissimilar target images. Our main contribution is to show how to formulate and solve this problem using a sequence of quadratic programs. We also show how to compute the layers so that simply changing the order in which they are stacked leads to different shadow images under a fixed lighting configuration. We produce physical prototypes, illustrating a variety of real and synthetic results using the sun as a light source, and validate our method with a range of experiments.

Layered attenuators can make appealing art pieces. They can be incorporated into personalized memorabilia, such as keyrings, because of their low manufacturing cost and simplicity. An intriguing application is using a layered attenuator in place of a stained glass window to produce different images depending on the day or the time of day. If only a fraction of layer permutations lead to images, finding them may be an amusing puzzle.

2. Related Work

Our work falls into the area of data-driven materials, a field that has recently received increased attention in computer graphics. An early example of such work is a method to generate bas-reliefs from 3D scenes [WDB*07]. More recently, several techniques detail how to construct physical objects, which manipulate incident light in order to produce a desired effect. Fuchs et al. [FRSL08] designed a light-dependent display using lenslet arrays for spatial multiplexing. Papas et al. [PJJ*11] presented a computational pipeline to manufacture custom smooth lenslet arrays, which produce a desired caustic image using specular refraction or reflection. Similarly, Weyrich et al. [WPMR09] design surface micro-structures to match a target BRDF or highlight profile. Band Moiré effects have also been used to produce images [HC04], where slight changes in alignment produce large changes in the apparent image. Our method similarly finds high-frequency patterns that exploit small changes in alignment due to different light directions to produce different images.

We create physical objects which cast desired colored shadow images when illuminated by pre-determined lighting configurations. This idea of exploiting light transport to generate multiple output images has been previously investigated. Alexa and Matusik [AM10] produce diffuse surfaces whose shading can appear as two distinct images when illuminated from different directions. This method is limited to grayscale images and tends to suffer from ghosting artifacts when the target images exhibit significant differences. Mitra and Pauly [MP09] designed and manufactured opaque 3D objects which, when lit from different directions, cast prescribed binary shadows. Both methods rely on idealized point or directional light sources.

As with our method, several recent papers have manufactured volumetric objects by stacking printed transparent layers. This approach provides a simple and convenient way of representing 3D data at the cost of quantizing depth. Holroyd et al. [HBLM11] build such layered objects, which, when viewed from a range of directions, give the illusion of an embedded 3D object. Wetzstein et al. [WLHR11] use printed multi-layer attenuators, coupled to a uniform backlight, to create hand-held light-field displays for static scenes. Displays capable of reproducing 4D light fields have also been manufactured by layering LCD panels [Got10, LWH*11]. This promising approach to manufacturing 3D displays circumvents the reduced spatial resolution of lenticular autostereoscopic displays [Lip08] and reduced brightness of parallax barrier displays [Ive03].

In order to cast targeted color shadows, our method reconstructs attenuation functions over a volumetric domain. In that respect, our work is similar in spirit to methods that use tomography for volumetric 3D reconstruction, a technique that is widely used in medicine and other fields [Her10]. In computer graphics, for instance, Ihrke and Magnor [IM04] compute the 3D structure of flames using sparse-view tomography. Trifonov et al. [TBH06] use visible-light tomography to solve for a volume of absorption values in order to reconstruct transparent objects. Atcheson et al. [AIH*08] reconstruct a volume of refractive indices to capture gaseous flow using bent-ray tomography. Wetzstein et al. [WLHR11] also use tomographic reconstruction to build the layered 3D display mentioned above.

These tomographic approaches rely on having an objective function that is quadratic in the layer pixels. Such an objective is usually obtained by formulating the problem in log space. Wetzstein [Wet11] shows that it is possible to obtain a solution using a general-purpose nonlinear solver without going to log space, but much more slowly. Working in log space, however, requires point or directional lighting as well as a pointwise error metric. We demonstrate in Section 4 that these constraints lead to heavy ghosting artifacts. We show that an area light formulation allows the optimization additional degrees of freedom and significantly reduces this ghosting. Using area lights with multiple attenuation layers results in a nonlinear image formation model, which cannot be linearized by moving to log space. Therefore, we solve for the attenuator layers by optimizing a sequence of quadratic programs, similarly to alternating least squares [Spe72] algorithms.

3. Approach

We express the shadow image for a given lighting configuration (or permutation) as a function of the attenuators. This function is linear in each attenuator individually, which enables us to write a quadratic program to optimize a single attenuator to match the target images as closely as possible in a least squares sense. We then iteratively optimize one attenuator at a time until convergence. Convergence is guaranteed since the error can never be below zero and never increases.

3.1. Shadow Image

We introduce our setup with a number of simplifying assumptions and later discuss which of them can be relaxed. We assume a white, diffuse receiver surface on the xy plane (Figure 2). We consider color channels independently. The equation for the irradiance at a point on the receiver is:

$$I(x,y) = \int_\Omega L_{\text{in}}(x,y,\omega)\, T_r(x,y,\omega)\, \max(0,\omega^z)\, d\omega, \qquad (1)$$

where L_{in} is the unimpeded incident radiance at (x,y) from direction ω, ω^z is the z component of ω, Ω is the domain of all unit direction vectors, and the transmittance function T_r denotes the attenuation that each ray of light undergoes as it passes through the multi-layer attenuator. We model the multi-layer attenuator as n infinitely thin layers, each parallel to the receiver at distances z_i, and with an image $A_i(x,y)$ that modulates light transmission. We can then write:

$$T_r(x,y,\omega) = \prod_{i=1}^{n} A_i(\text{hit}(x,y,\omega,z_i)), \qquad (2)$$

where $\text{hit}(x,y,\omega,z) = (x+\omega^x z/\omega^z, y+\omega^y z/\omega^z)$ is the ray-attenuator-layer intersection. We further assume our lighting is distant, so the unattenuated incident radiance can be written as an environment map: $L_{\text{in}}(x,y,\omega) = E(\omega)$. We estimate I by importance-sampling using the distribution $\text{pdf}(\omega) = E(\omega)\max(0,\omega^z)/c$, where $c = \int_\Omega E(\omega)\max(0,\omega^z)d\omega$. The estimate using m samples ω_j is:

$$\begin{aligned} I(x,y) &\approx \frac{1}{m}\sum_{j=1}^{m} \frac{L_{\text{in}}(x,y,\omega_j)\, T_r(x,y,\omega_j)\, \max(0,\omega_j^z)}{\text{pdf}(\omega_j)} \\ &= \frac{c}{m}\sum_{j=1}^{m}\prod_{i=1}^{n} A_i(\text{hit}(x,y,\omega_j,z_i)). \end{aligned} \qquad (3)$$

3.2. Optimization

We are given p target images $I_k(x,y)$ and lighting configurations $E_k(\omega)$. The distances of the attenuators to the receiver z_i are also prescribed. We importance-sample each $\text{pdf}_k(\omega)$ to obtain $\omega_{j,k}$ and optimize the total squared error:

$$\sum_{k=1}^{p}\sum_{x}\sum_{y}\left(I_k(x,y) - \frac{c_k}{m}\sum_{j=1}^{m}\prod_{i=1}^{n} A_i(\text{hit}(x,y,\omega_{j,k},z_i))\right)^2 \qquad (4)$$

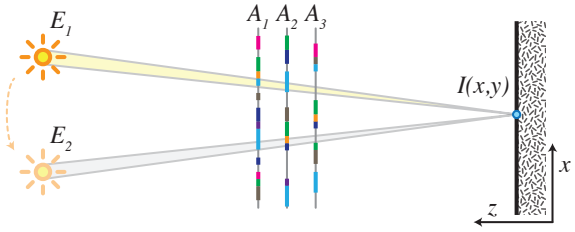

Figure 2: *Light rays from two different sun positions pass through different locations on the attenuators, allowing the formation of distinct images. The nonzero radius of the sun blurs the images, but provides flexibility in distributing error.*

subject to the constraint that $0 \le A_i(x,y) \le 1$. The standard approach to similar problems in tomographic reconstruction is to go to log space, where the product over the attenuators becomes a sum and the objective becomes quadratic. However, since we have an environment light, rather than a point or a directional light, this does not work because the logarithm of a sum does not generally simplify. Instead, we note that if we fix all but one of the attenuators A_i, the irradiance is linear in the remaining attenuator and therefore, the objective is convex quadratic. Representing the receiver irradiance as a $w_r \times h_r$ image $\mathbf{I}$ and each attenuator as a $w_a \times h_a$ image $\mathbf{A}_i$, and assuming no filtering, the irradiance is:

$$\mathbf{I} = \mathbf{M}_i\mathbf{A}_i, \qquad (5)$$

where $\mathbf{M}_i$ is a $w_r h_r \times w_a h_a$ matrix whose entry in row (x_r,y_r) on the reciever and column (x_a,y_a) on the attenuator is:

$$\frac{c}{m}\sum_{\substack{\omega_j \text{ from } (x_r,y_r)\\ \text{through } (x_a,y_a)}}\prod_{i'\neq i} A_{i'}(\text{hit}(x_r,y_r,\omega_j,z_{i'})). \qquad (6)$$

Written this way, the objective for attenuator i becomes:

$$\sum_{k=1}^{p}\left\|\mathbf{I}_k - \mathbf{M}_{i,k}\mathbf{A}_i\right\|^2. \qquad (7)$$

We optimize each attenuator in a round robin fashion using a sparse quadratic programming (QP) solver. Because each attenuator is optimized while keeping all the others fixed, the objective cannot decrease and the optimization converges. While we cannot guarantee convergence to a global minimum, our experiments show that results are not sensitive to initial conditions (see Figure 13).

3.3. Extensions

Although the implementation is simplest with the assumptions described above, the same basic algorithm can be used in more general settings. The lighting does not need to be distant: to support local lighting, we need a different set of importance samples for each (x,y) on the receiver and each light. To accommodate refracting physical spacers between the attenuators, the $\text{hit}(x,y,\omega,z)$ function can be adjusted to

compute the intersection of the refracted ray with the layer. Furthermore, the receiver does not need to be parallel to the attenuators or even planar and can vary with the lighting configuration—an appropriate hit function is the only change necessary for these cases.

The error metric we use above is pointwise mean squared error. Our formulation extends to alternative error metrics, better grounded in human perception: measuring the difference between the target image and the resulting irradiance under an arbitrary linear filter keeps the per-layer problem quadratic. This linear filter can, for example, be a blur kernel to enable the algorithm to distribute error better, a gradient filter to emphasize edges, or a combination of the two.

Instead of assuming square pixels, we can generalize the attenuators to be sums of basis functions $A_i(x,y) = \sum_b a_{i,b}\phi_b(x,y)$. In this case, in Equation (4), the product becomes a product-of-sums, which can be distributed into a sum-of-products and the objective remains quadratic in each attenuator separately. While basis functions can better model ink deposited by a printer, they increase the density of the matrices $\mathbf{M}_i$, which hurts performance. Because our extended light sources blur the pixels, we have not had to rely on general basis functions.

An interesting extension that we experimented with is to obtain different shadows by permuting the attenuators instead of changing the lighting configuration. In this case, k indexes a permutation and the matrices $\mathbf{M}_{i,k}$ are computed by keeping the lighting E fixed and changing the z_i's.

4. Results

We have implemented our algorithm in MATLAB using the built-in trust region sparse QP solver. To test our method, we used four sets of three images (Figures 3, 4, 5, 6). The attenuator layers were the same resolution as the target images, between 450 and 600 pixels on each side. We targeted a maximum image dimension of 15cm, for which this resolution was sufficient, but as Figure 14 shows, even 256×256 would have been sufficient for the attenuators. We used $m = 50$ samples for each light source.

We constructed physical prototypes, taking the sun at $+z$ and at $\pm 20^\circ$ to $+z$ as the three light sources. We printed the layers on transparencies with an off-the-shelf Epson Stylus Photo 1400 Inkjet. To improve contrast, we printed each layer on two transparencies and placed them together. We aligned the transparencies by printing small dots above the two top corners of each receiver and putting pins through them. We used 5mm thick acrylic spacers, but we do not account for the refraction of the acrylic in our computation; instead, we simply position the layered attenuator at a sharper angle to the sun, in accordance with Snell's law (about $\pm 30^\circ$). The receiver was assumed to be 30cm from the closest layer.

4.1. Experiments

Our results demonstrate that our method produces clear shadows even without precise alignment. Figures 1, 4, 5 and 7 (right) show the shadows cast by the layered attenuators we produced, each with three target images and three layers. We also produced attenuator layers that, when cyclically permuted, generate three images under the same lighting configuration (Figure 6). For a single layer, the irradiance on the receiver is essentially a convolution of the layer and E; computing even a single layer for a single target shadow is a deconvolution problem that cannot be solved exactly and some blur will inevitably be introduced.

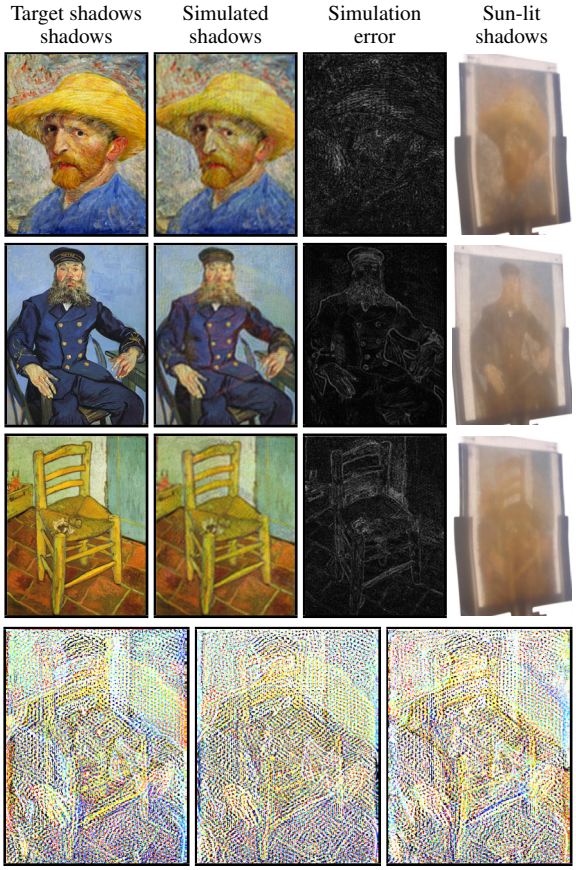

Figure 3: *Top: shadow images and error visualization for three Van Gogh paintings. Bottom row: the attenuator layers used to produce these results.*

To compare to standard tomographic approaches, we computed a layered attenuator assuming idealized (point and directional) light sources (Figure 7). Even when illuminated with the assumed ideal light in simulation, the result is considerably worse, showing heavy ghosting. To validate our optimization strategy in this case, we used log space code provided by Wetzstein et al. [WLHR11] and obtained qualitatively similar results with heavy ghosting. Wetzstein and

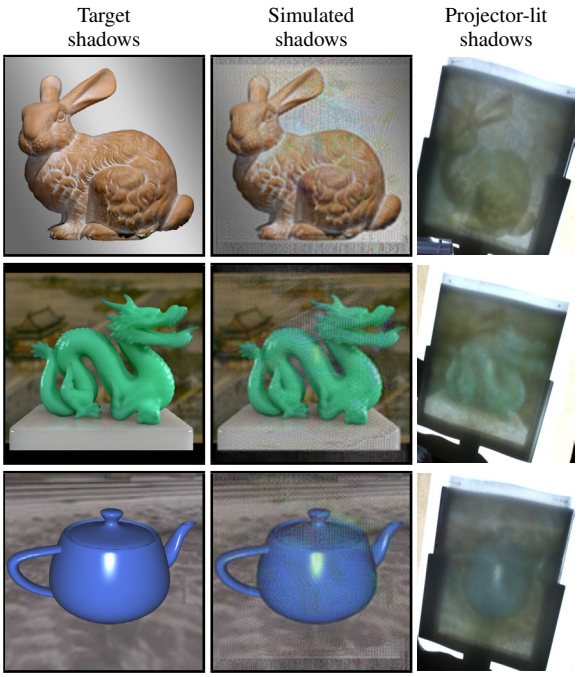

Figure 4: *Results using the usual suspects in computer graphics.*

Figure 5: *Results using vacation photos.*

Figure 6: *Shadow images generated by permuting the attenuator layers and keeping the sun light configuration fixed.*

colleagues discuss this effect as a limitation. This behavior is somewhat counterintuitive because, for one layer and one image, a directional light can produce a perfect image, while a light source like the sun inevitably introduces blur. For multiple target images, however, a different effect dominates: for a directional light, each pixel on the receiver is only affected by a small number of points on the attenuators; in this setup, there are not enough degrees of freedom to distribute error spatially. In contrast, for an extended light source, like the sun, each receiver point is influenced by many points on the attenuator. The blur actually helps mask the error in this case (cf. blurring a dithered image). In our case, the cone from a point on the reciever to the sun has a radius of about 1.4mm on an attenuator 30cm away from the receiver. This corresponds to just under 5 pixels at our resolution. The area of this circle is over 70 pixels, providing sufficiently many degrees of freedom for our optimization.

The simulated results we show look much better than the physical results and we attribute this to several factors. Our pipeline assumes a perfectly white sun and an RGB light model, but the reality is that the spectra of the sun, the printer inks, and the observer (eye or camera) have a nontrivial interaction. Alignment errors introduced when the layers are stacked can also significantly degrade the result quality. There may be scattering in the acrylic, transparencies, and printer ink that further adds to the blur. Moreover, the simulation error is not always a good predictor of how good a result will look when manufactured. For example, we opti-

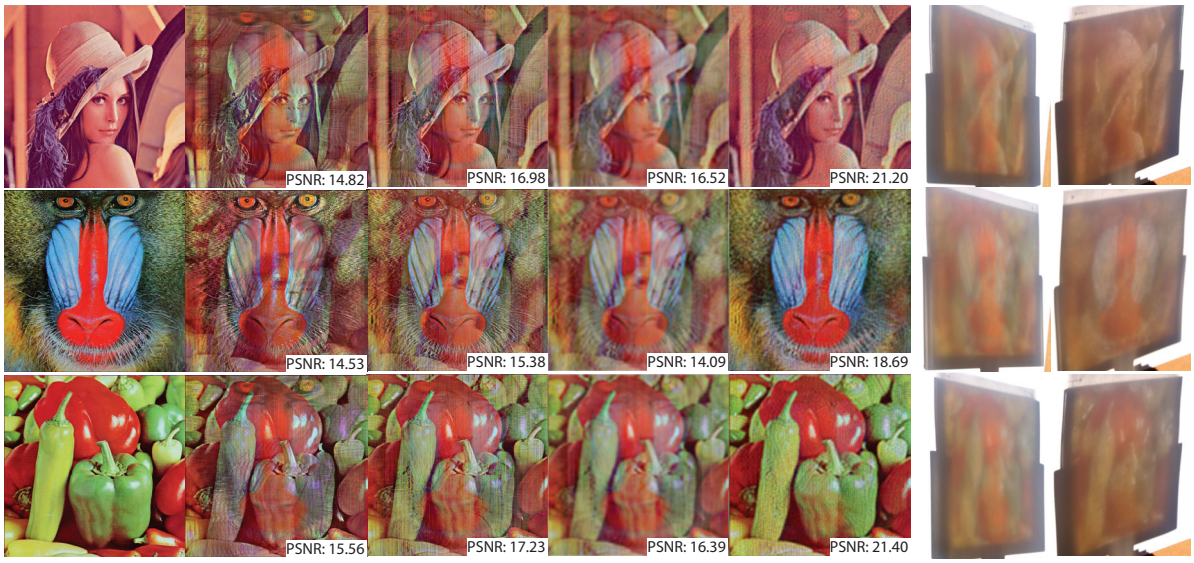

Figure 7: *Comparison of extended vs. idealized light sources. (a) Target images. (b) Ideal simulation of layers optimized for directional light in log space [WLHR11]. The PSNR is lower than in column c because the error is not measured in log space. A directional light is used instead of a point light because that is what their implementation supports. (c) Ideal simulation of result optimized for a point light using our round-robin optimization. (d) The same result simulated using the sun, an extended light soure. (e) Our result, optimized and simulated using the sun. (f) The physical shadow produced by the layers in column c. (g) The physical shadow produced by our result.*

mized a four-layer attenuator for nine images and nine different sun positions. Although the synthetic reconstruction is surprisingly good (Figure 8), when we manufactured this result, we were only able to see a blur.

We have experimented with varying the number of layers, while keeping the same three target images (Figure 9). Somewhat surprisingly, even a two-layer attenuator can produce three target images whose quality is not much worse, even in reality (Figure 10), than for a three-layer attenuator (Figure 3). In practice, this may be partly due to the fact that aligning more layers is harder, and that slight changes to sun position can compensate for misalignments between two layers.

Our implementation is not optimized, but since the runtime is dominated by the QP solver (for 512×512 attenuator layers), there is not a lot of room for improvement without improving the solver. When solving for an attenuator layer image, we pass the current version of the layer to the solver as an initial guess and set a relatively high termination tolerance, so that the solver does not spend too long finding the optimal solution to a QP that will be different in the next iteration. Computing a color, three-image three-layer attenuator with 512×512 images and attenuators takes just over two hours on a Core i7 970 with 12GB of RAM.

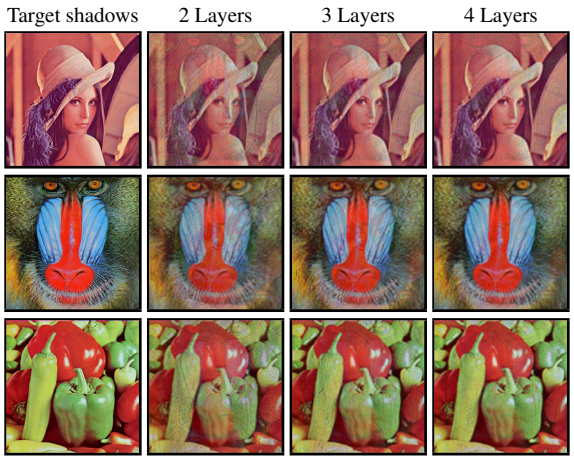

Figure 9: *Simulated results with two–four attenuator layers.*

4.2. Convergence and Error

We performed several quantitative experiments to measure the quality of the results produced by our optimization. We measure error by taking the peak signal-to-noise ratio (PSNR) with respect to the mean squared error averaged over all images and color channels. Figure 11 shows how our optimization converges for different numbers of layers. One iteration corresponds to n QP solutions, one for every layer. In Figure 12 we explore how error increases as we increase

Figure 8: *Even with nine target images (top), a four-layer attenuator generates high-quality simulated results (bottom). However, when fabricated, the result is unusable.*

Figure 10: *A layered attenuator with only two layers produces three distinct shadows.*

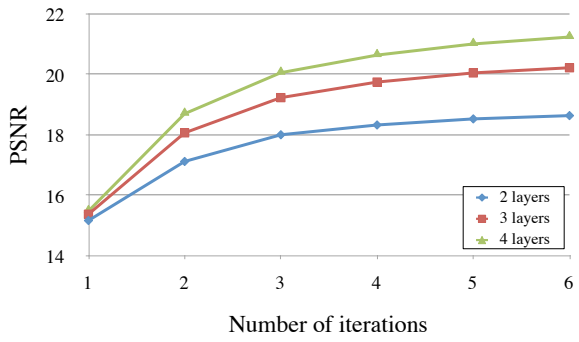

Figure 11: *Average error of the Lena-Mandrill-Peppers test as our optimization progresses for attenuators with different numbers of layers. Figure 9 shows the final results.*

the number of images. Figure 13 demonstrates our robustness to initial conditions: the performance of the algorithm is pretty much the same no matter what attenuator layers we start with.

5. Conclusion

We introduce a framework for controlling the appearance of shadows cast through volumetric attenuators. We manufacture multi-layer attenuators by stacking a set of commodity transparencies with printed patterns automatically generated by our framework. We can simultaneously encode several

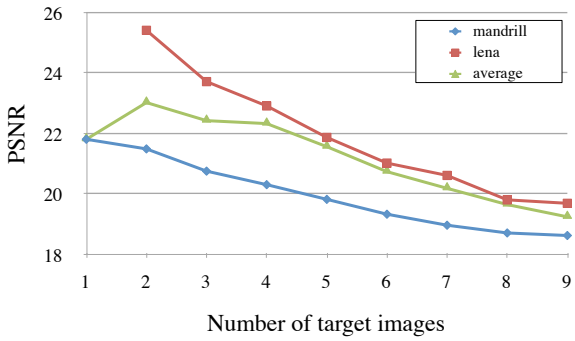

Figure 12: *Algorithm stress-test: we compute four-layer 256×256 grayscale attenuators for increasing numbers of images, up to nine. We plot how the error of the first two images increases as more images were added, as well as how average error changes. In this test, average error decreases as the second image is added because it has fewer high frequencies and is easier to reproduce than the first image.*

prescribed color shadows, which become visible as the incident light changes. Alternatively, we can control the appearance of the cast shadows by simply permuting the stacked layers, while keeping the lighting fixed.

Limitations. Several effects are not currently taken into account by our simulation. Calibrating for the printer inks, the light source, transparency attenuation and ambient light, and using spectral light transport should decrease the gap between real and simulated results significantly (albeit at the expense of a more complicated pipeline). More precise layer alignment, and modeling the refraction and reflection in the layered attenuators should further decrease the real error.

Our method works well with concentrated light sources, such as the sun. More diffuse sources, such as full environment maps, are more challenging since shadow intensities tend to be weaker. Using a gradient domain error measurement may nevertheless produce recognizable images in that case. Finally, our long running times prohibit interactive layered attenuator design; better initial guesses and specialized optimization algorithms could speed up the process.

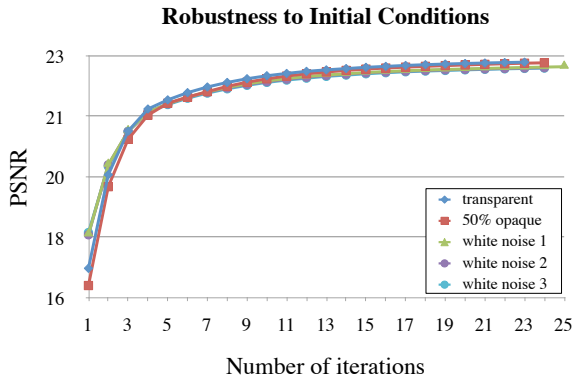

Figure 13: *Plot of the three-layer Lena-Mandrill-Peppers optimization progress, depending on how layers were initialized. There is essentially no difference between using fully transparent, partially opaque, or random initializations.*

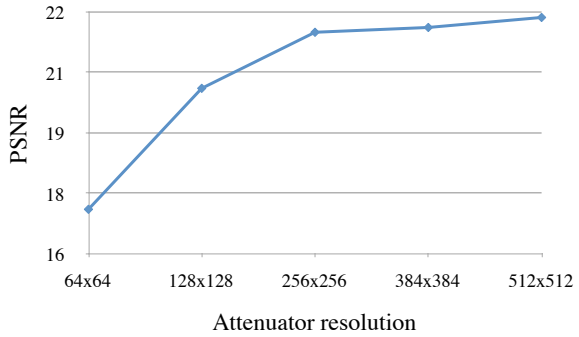

Figure 14: *For the three-layer Lena-Mandrill-Peppers test, we compare attenuator resolutions with a fixed target resolution of 512×512. The blurring effect of the sun allows 256×256 attenuators to be almost as accurate as 512×512.*

Future Work. An interesting avenue of future work is to model the effects of non-diffuse reflection on the receiver surface, for example, using a basis expansion of the outgoing radiance at each (x,y) location. While doing so introduces additional optimization parameters (and an extra sum to Equation 4), one could imagine generating shadow images that not only vary with the incident lighting configuration, but also with the position of the viewer.

Acknowledgments

We thank Aaron Hertzmann and Michael Daum for initial motivating discussions, and Gordon Wetzstein for helpful advice regarding the printing process.

References

[AIH*08] ATCHESON B., IHRKE I., HEIDRICH W., TEVS A., BRADLEY D., MAGNOR M., SEIDEL H.-P.: Time-resolved 3d capture of non-stationary gas flows. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 27*, 5 (2008), 132. 2

[AM10] ALEXA M., MATUSIK W.: Reliefs as images. *ACM Transactions on Graphics 29*, 4 (July 2010), 60:1–60:7. 2

[FRSL08] FUCHS M., RASKAR R., SEIDEL H.-P., LENSCH H. P. A.: Towards passive 6d reflectance field displays. *ACM Transactions on Graphics 27*, 3 (Aug. 2008), 58:1–58:8. 2

[Got10] GOTODA H.: A multilayer liquid crystal display for autostereoscopic 3d viewing. *Stereoscopic Displays and Applications XXI 7524* (2010), 1–8. 2

[HBLM11] HOLROYD M., BARAN I., LAWRENCE J., MATUSIK W.: Computing and fabricating multilayer models. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 30*, 6 (Dec. 2011). 2

[HC04] HERSCH R. D., CHOSSON S.: Band moiré images. *ACM Transactions on Graphics 23*, 3 (Aug. 2004), 239–247. 2

[Her10] HERMAN G. T.: *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*. Advances in Pattern Recognition. Springer, 2010. 2

[IM04] IHRKE I., MAGNOR M.: Image-based tomographic reconstruction of flames. In *ACM Siggraph / Eurographics Symposium Proceedings, Symposium on Computer Animation* (June 2004), pp. 367–375. 2

[Ive03] IVES F. E.: Parallax stereogram and process of making same. *U.S. Patent 725,567* (1903). 2

[Lip08] LIPPMANN J. F. G.: Epreuves reversibles donnant la sensation du relief. *Journal of Physics* (1908). 2

[LWH*11] LANMAN D., WETZSTEIN G., HIRSCH M., RASKAR R., HEIDRICH W.: Polarization fields: Dynamic light field display using multi-layer LCDs. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 30*, 6 (Dec. 2011). 2

[MP09] MITRA N. J., PAULY M.: Shadow art. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 28*, 5 (2009). 2

[PJJ*11] PAPAS M., JAROSZ W., JAKOB W., RUSINKIEWICZ S., MATUSIK W., WEYRICH T.: Goal-based caustics. *Computer Graphics Forum (Proceedings of Eurographics) 30*, 2 (June 2011), 503–511. 2

[Spe72] SPENCE I.: A monte carlo evaluation of three nonmetric multidimensional scaling algorithms. *Psychometrika 37*, 4 (1972), 461–486. 2

[TBH06] TRIFONOV B., BRADLEY D., HEIDRICH W.: Tomographic reconstruction of transparent objects. In *Proceedings of Eurographics Symposium on Rendering* (2006), pp. 51–60. 2

[WDB*07] WEYRICH T., DENG J., BARNES C., RUSINKIEWICZ S., FINKELSTEIN A.: Digital bas-relief from 3d scenes. *ACM Transactions on Graphics 26*, 3 (July 2007), 32:1–32:7. 2

[Wet11] WETZSTEIN G.: *Computational Plenoptic Image Acquisition and Display*. PhD thesis, Department of Computer Science, University of British Columbia, 2011. 2

[WLHR11] WETZSTEIN G., LANMAN D., HEIDRICH W., RASKAR R.: Layered 3D: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Transactions Graphics (Proceedings of SIGGRAPH) 30*, 4 (2011). 2, 4, 6

[WPMR09] WEYRICH T., PEERS P., MATUSIK W., RUSINKIEWICZ S.: Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (Proceedings of SIGGRAPH) 28*, 3 (July 2009), 32:1–32:6. 2