Printed Perforated Lampshades for Continuous Projective Images

HAISEN ZHAO, LIN LU and YUAN WEI
Shandong University
DANI LISCHINSKI
The Hebrew University of Jerusalem
ANDREI SHARF
Ben-Gurion University
DANIEL COHEN-OR
Tel Aviv University
BAOQUAN CHEN
Shandong University

We present a technique for designing 3D-printed perforated lampshades that project continuous grayscale images onto the surrounding walls. Given the geometry of the lampshade and a target grayscale image, our method computes a distribution of tiny holes over the shell, such that the combined footprints of the light emanating through the holes form the target image on a nearby diffuse surface. Our objective is to approximate the continuous tones and the spatial detail of the target image, to the extent possible within the constraints of the fabrication process.

To ensure structural integrity, there are lower bounds on the thickness of the shell, the radii of the holes, and the minimal distances between adjacent holes. Thus, the holes are realized as thin tubes distributed over the lampshade surface. The amount of light passing through a single tube may be controlled by the tube's radius and by its orientation (tilt angle). The core of our technique thus consists of determining a suitable configuration of the tubes: their distribution across the relevant portion of the lampshade, as well as the parameters (radius, tilt angle) of each tube. This is achieved by computing a capacity-constrained Voronoi tessellation over a suitably defined density function, and embedding a tube inside the maximal inscribed circle of each tessellation cell.

Categories and Subject Descriptors: I.3.3 [**Computer Graphics**]: Picture/Image Generation—*Display algorithms*; I.3.8 [**Computer Graphics**]: Applications

Additional Key Words and Phrases: Light projection, perforation

1. INTRODUCTION

The emergence of 3D printing technology opens new interesting opportunities for computerized halftoning. In this paper, we introduce the generation of 3D-printed perforated lampshades designed to project grayscale images onto surrounding surfaces. The spatial distribution of tone across the projection is generated by carefully controlling the amount of light shining through the printed surface. Similarly to the dots used in halftoning, our basic idea is to perforate the physical 3D shell of the lampshade with small tubes, whose parameters are used to control the amount and the spatial distribution of light that emanates from the lamp and reaches a receiving surface. Figure 1 shows an example of such a lampshade and its projected image. Projecting continuous tone imagery via a perforated surface can thus be regarded as *halftoning with light* or *3D halftoning*.

Halftoning is a technique used to represent continuous shades of gray through the use of discrete dots of the same color, varying either in size, in shape, or in spacing. The method relies on the natural low-pass spatial filtering of the human visual system that blends the discrete dot pattern into a continuous tone. Traditional halftoning uses simple-shaped dots, typically circular, of continuously varying size. However, when the dots all have the same size, spatial resolution can be traded for perceived tone resolution [Ulichney 1988; Lau and Arce 2008].

The 3D halftoning technique that we present differs from ordinary halftoning in a number of ways. First, unlike other common digital media, here one can generate dots of continuous sizes. In that sense, the technique is closer to analog halftoning, where the dot sizes are continuous. Second, dots are holes in a surface, realized as tiny 3D tubes, hence having both *radius* and *orientation* with respect to the light source. Third, the resulting image is not printed but projected, which requires to consider the geometry of both the projecting surface (the lampshade) and the receiving surface (the wall). Lastly, since the projected light footprint of each tube is slightly blurred, and multiple footprints add up in areas of overlap, there is already some low-pass filtering inherent in the image formation process. We explicitly account for the effect of overlapping footprints of adjacent tubes, which is inhibited in digital halftoning, and take advantage of it.

Moreover, it should be stressed that in our setting, the major challenge is to ensure the printability and structural integrity of the perforated lampshade. Specifically, there are strict lower bounds on the radii of the tubes, and on the inter-tube spacing. Violating the former constraint would result in clogged tubes, and the latter would render the shell fragile and prone to breakage. In particular, having a lower bound on the tube radius means that darker tones cannot

Corresponding authors: Lin Lu (lulin.linda@gmail.com), Baoquan Chen (baoquan.chen@gmail.com)

DOI: 10.1145/XXXXXXX.YYYYYYY

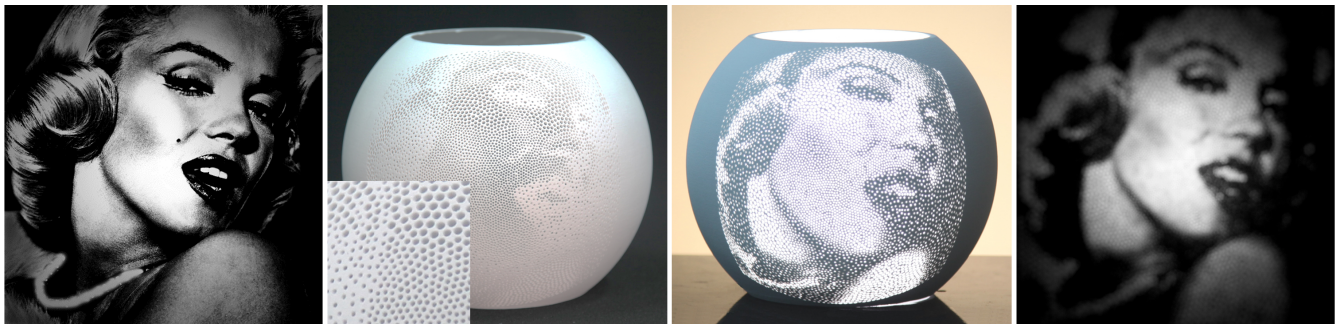

Fig. 1. A 3D-printed lampshade projecting a continuous grayscale image of Marilyn Monroe. From left to right: the target image, the 3D-printed perforated lampshade, the lampshade with a light source inside, and the resulting projection onto a nearby wall. The superposition of the individual light footprints on the wall forms a visually continuous image.

be achieved by simply using smaller holes; instead, we reduce the amount of light passing through a tube by tilting it away from the light source center.

We present a 3D halftoning technique that, given the geometry of the lampshade surface, a target grayscale image, and a receiving surface, produces a spatial distribution of tubes, along with their radii and orientations, such that the resulting projected image faithfully reproduces the input image. The goal is to produce a continuous tone projected image, which strives to match the distribution of tones and the spatial detail of the original.

The target intensity at each projected location is matched by placing tubes with suitable radii and tilt angles around the corresponding location on the lampshade. Brighter areas are reproduced using tubes with a larger radius, while in darker areas we place minimal radius tubes, tilted away from the light source center. In order to maximize the spatial resolution, the tubes must be placed as densely as possible, but they must not violate the inter-tube distance constraint. To achieve this, we embed each tube inside a disk that incorporates a safety margin around the tube. Note that differently from ordinary halftoning or stippling, in our case both brighter and darker areas require larger disks, with the maximal density of disks corresponding to the middle of the grayscale range.

Having reduced the problem to one of finding a dense packing of disks with spatially varying radii, we solve it by computing a capacity-constrained Voronoi diagram over a suitably defined density function, whose value at each location is inversely proportional to the required disk area. The process is illustrated in Figure 2.

In summary, the contributions of this paper are as follows:

—We tackle the novel problem of generating projected imagery by shining light through a 3D-printed perforated surface.

—We introduce a novel 3D halftoning approach, where the halftoning dots are realized as a distribution of tubes passing through a solid surface, with varying radii and orientations.

—We present a method that determines the spatial distribution and the parameters of the individual tubes, while striving to match the target image, subject to fabrication constraints.

Our results demonstrate the effectiveness of the proposed technique using a variety of target images. The limitations of the process are also demonstrated and discussed.

2. RELATED WORK

2.1 Halftoning and Stippling

Halftoning is a classical technique which played a major role in traditional paper printing and in digital displays [Kipphan 2001; Ulichney 1987]. Our work leverages this classical technique in the new domain of digital fabrication and 3D printing.

Throughout the evolution of digital halftoning over the past decades, the dominating issues have persistently been ones point density and spatial resolution, reduction of noticeable regular patterns [Mitsa and Parker 1992], and preservation of structural details [Eschbach and Knox 1991; Pang et al. 2008; Chang et al. 2009; Li et al. 2010].

Similarly to halftoniong, stippling also uses spatially-varying dot patterns to convey shading. Balzer et al. [2009] propose using the capacity-constrained Voronoi tessellation (CCVT) for enforcing the constraint of equal weighted area for the region around each point in a stippled image. However, the method they propose suffers from high computational complexity. Consequently, considerable attention was given to developing faster alternatives. In particular, de Goes et al. [2012] formulate CCVT as a constrained optimal transport problem. This results in a fast continuous constrained minimization method, which is able to enforce the capacity constraints exactly. In this work, we show how to encode the objectives and the requirements of our 3D halftoning method as a density function, which enables solving the problem by constructing a CCVT, and use the method of de Goes et al. for this purpose.

The works mentioned earlier are concerned with generating halftoning and stippling patterns in 2D, while our work focuses on creating a perforation pattern on a 3D surface with finite thickness, in the context of 3D printing. In this context, Stucki [1997] first introduced the idea of 3D digital halftoning for transforming continuous-density objects into binary representations for rendition with additive fabrication technologies. Lou and Stucki [1998] further adapt the ordered-dither and error diffusion algorithms to the 3D case, focusing on stereolithography (SLA) printing techniques. Zhou and Chen [2009] utilize 3D digital halftoning by replacing a dot with a droplet, to achieve satisfied approximation accuracy with larger layer thickness, thus reducing the fabrication time for layered based inkjet printing processes. In contrast to these works, our 3D halftoning method is designed for the purpose of forming projected imagery that reproduces the tones of a target image, subject to fabrication constraints.

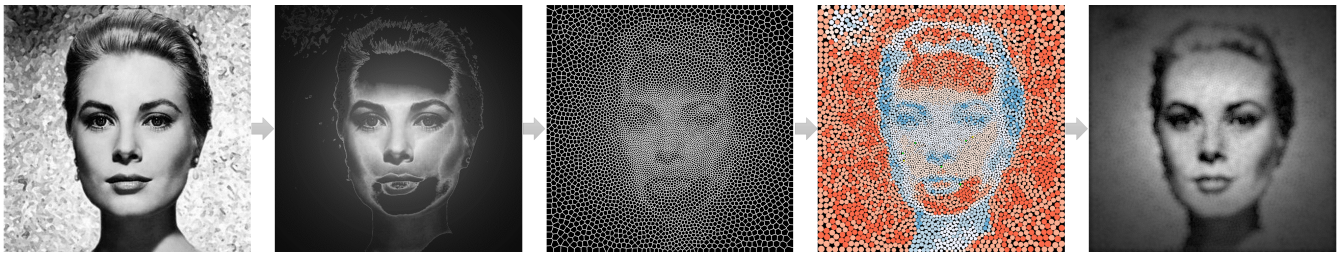

(a) (b) (c) (d) (e)

Fig. 2. Process overview: given a target image (a) we compute a density map (b) whose value at each location is inversely proportional to the required disk area (as explained in Section 3.1). Next, we compute a capacity-constrained Voronoi tessellation (c), and inscribe a maximal disk inside each cell (d). Red shades indicate disks where the radius of the embedded tube is larger than the minimal size (increasing the brightness), while blue shades indicate disks where the embedded tube is tilted (decreasing the brightness). The simulated projected image is shown in (e).

2.2 Optics-related Fabrication

Recently, design or generation of illumination effects via geometric modulation has been drawing increasing attention. Mitra and Pauly [2009] introduce shadow art, an algorithm for computing a 3D volume, whose shadows best approximate multiple binary images. Subsequent approaches generalize this to colored shadows using volumetric objects manufactured by transparent acrylic [Wetzstein et al. 2011; Baran et al. 2012]. In general, they build transmittance functions to simulate light attenuation through multi-layered attenuators.

Our work is inspired by the technique of Alexa and Matusik [2012], who drill holes with varying depths on a surface to induce the given image based on the occlusion of small holes.

Research in this area has also addressed controlling light reflections and refractions using surface modeling methods. These works cast a desired caustic image by modulating the geometry using microfacets [Weyrich et al. 2009], micropatches [Papas et al. 2011], B-spline surfaces [Finckh et al. 2010], continuous surfaces [Kiser et al. 2013; Yue et al. 2014] or normal fields [Schwartzburg et al. 2014], and then milling or engraving the surfaces. Researchers have also shown light effects in a wide range of applications, such as steganography [Papas et al. 2012], image display [Malzbender et al. 2012; Levin et al. 2013] and appearance design [Lan et al. 2013].

Printing optical fibers that control light propagation through total internal reflection between two surfaces has drawn some research attention thanks to the cutting edge multi-material 3D printers. Willis et al. [2012] design optical fibers for customizing interactive devices in display and illumination. Pereira et al. [2014] introduce automatic fiber routing algorithm that minimizes curvature and compression for optimal light transmission.

All of the above are high-end techniques, which rely on expensive manufacturing equipment and specific materials. In contrast, our method is designed for standard 3D printing technologies and common, widely used fabrication materials, such as plastic and powder. For such materials, the light transmittance is very low, as light mostly scatters due to the surface roughness. Therefore, shape modulation for light or shadow effects based on surface reflection and refraction cannot be easily adapted. In particular, using varying material thickness to control the projected image intensity is infeasible, and our approach is designed for lampshade shells of uniform thickness.

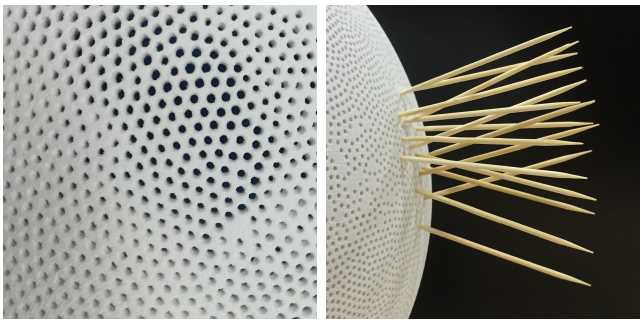

(a) (b)

Fig. 3. The physical 3D-printed lampshade shell perforated with tubes of varying sizes and density (a), and different tilt angles (b), as may be seen from the directions of the toothpicks.

3. PERFORATED LAMPSHADE DESIGN

Recall that our goal is to design a 3D-printable perforated lampshade, such that the light emanating from the lamp forms a continuous tone image, which is as close as possible to a target grayscale input. Light emanates from the lamp through holes in the lamp's surface. Due to the finite thickness of the lamp's shell, the holes may be thought of as tiny tubes, whose density, radii, and orientations vary across the surface, as may be seen in Figure 3. Since the total distribution of light across the projected image is determined by the combined effect of light passing through the tubes, our challenge is to determine the tubes' parameters, while respecting fabrication constraints imposed by the need to obtain a printable and structurally sound surface.

Specifically, given a target grayscale image I^t our task is to configure a set of tubes perforating the lampshade shell, such that the following objectives and requirements are satisfied:

a) The projected image I^p on a given diffuse surface, henceforth referred to as the *wall*, closely approximates the target grayscale values across the projection. Note that the approximation is up to some scaling factor, since we have control over the total amount of luminous flux emitted by the light source, as well as over the exposure time, when capturing the projection with a camera.

b) The projected image should exhibit continuous tones, while resolving fine spatial detail, as much as possible. In order to

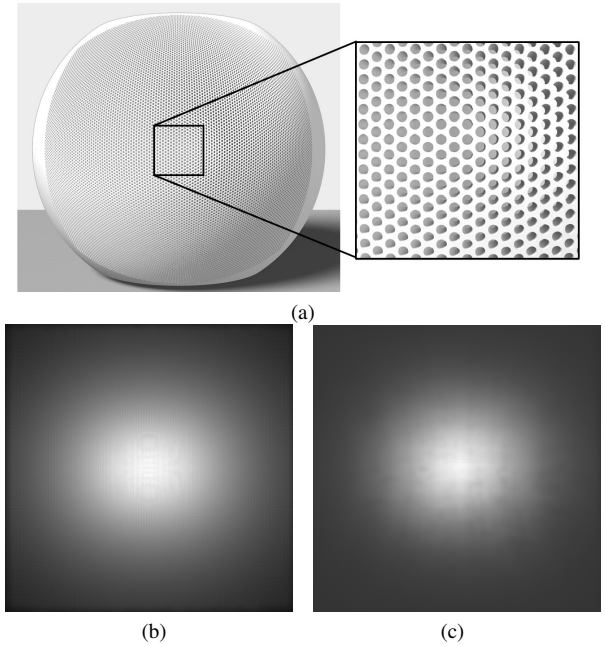

Fig. 4. (a) Maximally dense packing of the smallest printable tubes, oriented towards the center of the light source. (b) The corresponding simulation result. (c) A photograph of the actual projected image corresponding to this pattern.

achieve this objective the density of the tubes should be maximized, while their radii should be minimized.

c) Fabrication constraints: (a) the radius of a tube is bounded below by $r_{\min}$ to prevent clogging during the manufacturing process; and (b) any two adjacent tubes must have a gap of width greater than $d_{\min}$ of solid material between them. When the latter constraint is violated, the shell becomes too fragile.

The densest arrangement of tubes that satisfies the fabrication constraints above is shown in Figure 4(a). This arrangement corresponds to a hexagonal packing of disks of radius $r_{\min} + 0.5\, d_{\min}$, with a tube of radius $r_{\min}$ embedded at the center of each disk. The resulting projected pattern is shown in Figure 4(c). Our setup uses a disk-shaped approximately Lambertian light source that faces a planar projection surface. Thus, the intensity falls off away from the center due to the increase in distance and in the angles between the outgoing and incident light to the surface normals of the source and the receiving surface, respectively. In order to match a target image I^t, the tube arrangement must be adjusted in a spatially varying fashion. In order to achieve brighter tones we must increase the tube radii, but darker tones cannot be achieved by decreasing the radii, since this would violate the first fabrication constraint above.

One alternative is to employ a stippling approach to generate a variable density set of tube positions that would satisfy the fabrication constraints above. However, such an approach necessitates spacing tubes much farther apart from each other in the darker areas of the image, resulting in projection of isolated light dots inside such areas. Thus, continuity and spatial detail resolution has to be severely compromised. Instead, we propose decreasing the amount of light passing through tubes in darker areas by tilting their directions away from the light source center. This solution avoids unnecessarily isolating tubes, and plays an essential role in achieving continuous gray tones in the projected imagery. It is worth mentioning that the tilting angle is not restricted by fabrication capabilities, which offers enough variable tuning range.

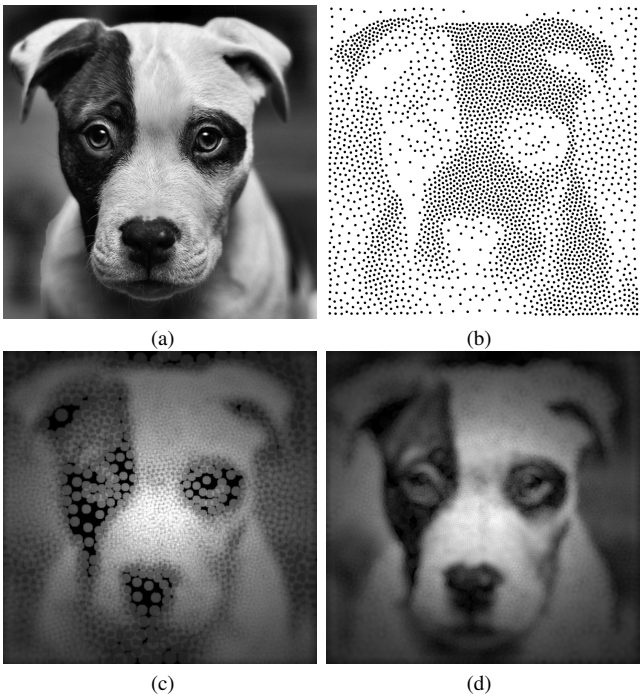

Fig. 5. For an input image (a), if we directly apply the stippling result (b) while ensuring the printability constraints, continuity and target grayscale tones cannot be achieved simultaneously in the projected imagery (c). Our method achieves a continuous result (d) by employing tube tilting to darken the tone. Here, (c) and (d) are both simulated images.

The inadequacy of a stippling-based approach is demonstrated in Figure 5. Here, starting from an input image (a), we use a state-of-the-art stippling method [de Goes et al. 2012] to generate the pattern in (b). In order to satisfy the fabrication constraints, the number of dots in this pattern is limited to roughly 3K. This results in a (simulated) projected image with poor continuity and isolated light spots, especially inside the darker tone areas around the dog's eyes and nose. Using our proposed solution, we are able to distribute roughly 6K tubes across the lampshade, thereby achieving the same apparent grayscale range, but in a much more continuous fashion.

Note that in order to ensure a minimal inter-tube distance, both widened tubes and tilted tubes must be placed inside larger disks on the lampshade surface, as illustrated in Figure 8. The radius of each disk is chosen so as to accommodate the widened or tilted tube, in addition to a *safety margin* of width $0.5\, d_{\min}$. In summary, each target intensity at a particular location corresponds to a certain disk radius, and our goal is to distribute a set of tightly packed non-overlapping disks with the desired radii across the relevant portion of the lampshade surface.

We attempt to achieve the required disk arrangement by computing a capacity-constrained Voronoi tessellation over a suitably defined density function, and inscribing a maximal circle inside each tessellation cell. The density function for a particular target image is determined from the differences between the target image intensity and the simulated projected image corresponding to the maximum uniform density pattern, shown in Figure 4(b).

The entire design process is visualized in Figure 2 and summarized in Algorithm 1. The main stages of the process are described in more detail in the remainder of this section.

Algorithm 1 Perforated Lampshade Design

Input: Grayscale image I^t, lampshade surface S_L, receiving wall surface S_W, light source
Output: Set of tube parameters (position, radius, tilt) on S_L

Simulate the reference image

—Simulate the projected *reference image* B_0, formed by tubes of radius $r_{\min}$, oriented at the light source center. Each tube is embedded in a disk of radius $r_{\min} + 0.5d_{\min}$, with the disks arranged in a tight hexagonal grid on the relevant portion of S_L (see Figure 4).

Compute the density function on S_W

—For each image pixel at position (x, y), determine the desired disk radius r_l on S_L following Equations (1-3).
—Compute the corresponding disk radius r_w on S_W.
—Set the density $\rho_w(x, y) \propto 1/r_w^2$, and normalize.

Compute the disk distribution

—Determine the desired number of disks N.
—Compute a capacity-constrained Voronoi tessellation using N disks and the density given by ρ_w.
—Inscribe the maximal circle inside each tessellation cell.

Generate tubes

—For each disk of radius r, centered at (x, y) on S_W: If $I^t(x, y) \geq B_0(x, y)$, the tube radius is set to $r - 0.5d_{\min}$ (the widest tube that fits inside the disk, with a safety margin, see Figure 8). The tube is oriented towards the center of the light source.
—If $I^t(x, y) < B_0(x, y)$, the tube remains with radius $r_{\min}$, and tilted away from the light source center (in a random direction), by the maximal angle afforded by a disk of radius $r - 0.5d_{\min}$ (see Figure 8).

3.1 Computing the target density function

Given the lamp setup (lampshade and wall geometry), we compute the densest arrangement of tubes that satisfies the fabrication constraints. Thus, the surface is tessellated by uniform hexagonal cells, and for each cell the maximal inscribed circle is with radius $r_{\min} + 0.5d_{\min}$, such that the smallest tube (radius $r_{\min}$) can be perforated. We take the projected image as the reference B_0. The arrangement, the result of the simulation, and a photograph of the actual projected image are shown in Figure 4. The simulation process is described in detail in Section 3.4.

Image B_0 indicates the darkest continuous tones that can be achieved via varying radii; hence, darker tones are achieved by tilting tubes. Given a target image I^t, the reference B_0 enables us to determine the tube mode, radius enlarging or tilting. For both modes, an embedding disk with radius larger than $r_{\min} + 0.5d_{\min}$ is needed for shape modulation. Then we will define a density function on the wall based on the desired radii, which is used to compute a capacity-constrained Voronoi tessellation.

Since the projected image is formed on the wall by a linear combination of light transmitted through the tubes, we first linearize the given target image I^t by applying an inverse gamma correction. We assume that the target image is gamma corrected with a standard gamma value of 2.2.

For each projected image pixel (x, y), whose target grayscale value is given by $I^t(x, y)$, we match the target intensity by controlling the degree of occlusion of the light source by the perforated lampshade:

$$I^t(x, y) = \mathcal{K} E_v^0(x, y), \tag{1}$$

where $E_v^0(x, y)$ is the unoccluded illuminance of (x, y) on the wall (eq. 4), and $\mathcal{K}$ indicates the proportion of light unoccluded by the lampshade. For a uniform hexagonal tessellation on the lampshade, inside each cell we have $\mathcal{K} = Area(\text{unoccluded})/Area(\text{Cell})$. Approximating the area of each cell by that of a planar hexagon, we can express $\mathcal{K}$ as a function of r.

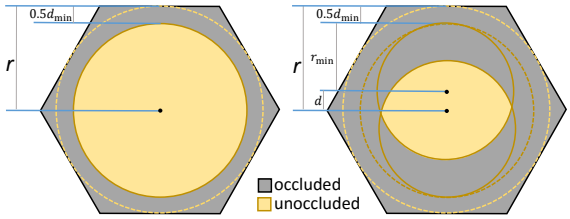

Fig. 6. The unoccluded area inside one hexagonal cell on the lampshade. Left: untilted tube case; Right: tilted tube case.

Specifically, for the untilted case (Figure 6 left), let r be the radius of the maximal inscribed circle inside the hexagon, then the unoccluded area is a disk of radius $r - 0.5d_{\min}$ and we have:

$$\mathcal{K}(r) = \frac{\pi(r - 0.5d_{\min})^2}{2\sqrt{3}r^2}. \tag{2}$$

For the tilted case, the tilted tube has the radius $r_{\min}$, and it is tilted by the maximal angle afforded by a disk of radius $r - 0.5d_{\min}$. The unoccluded area is the intersection of two disks (see Figure 6 right and Figure 8), and we have:

$$\mathcal{K}(r) = \frac{r_{\min}^2 cos^{-1}(d/r_{\min}) - d\sqrt{r_{\min}^2 - d^2}}{\sqrt{3}r^2}, \tag{3}$$

where $d = r - r_{\min} - 0.5d_{\min}$.

Thus, starting with the densest manufacturable arrangement of tubes, by comparing the corresponding projected image B_0 to the target image I^t we determine the locations where the tubes must be enlarged or tilted. Specifically, if $B_0(x, y) \leq I^t(x, y)$ the relevant tubes must be enlarged with the desired cell radius determined by equation (2). On the other hand, if $B_0(x, y) > I^t(x, y)$, the relevant tubes must be tilted and the desired cell radius is determined by equation (3).

In our current implementation, we compute the capacity-constrained Voronoi tessellation on the planar wall domain. Thus, we need to convert the desired radii on the lampshade to ones on the wall. Specifically, given a radius r of a disk on the lampshade, we compute the radius r_w of a disk on the wall, such that its projection back onto the lampshade (ellipse in the spherical case) has the same area as the original disk of radius r.

Given the target disk radii r_w across the projection area, we set the density to be proportional to the inverse of the disk area,

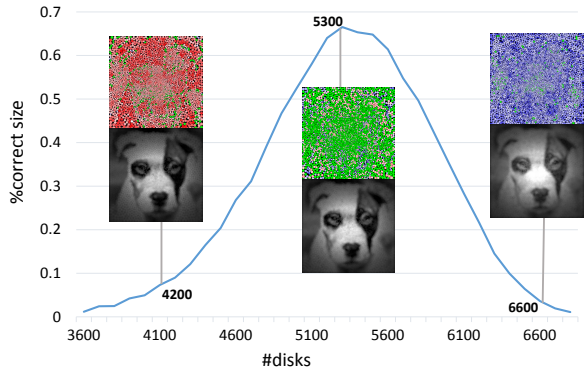

Fig. 7. The percentage of correctly sized disks is plotted as a function of the total number of disks. The disks and corresponding simulated images are shown for three locations on this curve. Disks that achieve their intended radius are indicated by green shades. Red and blue disks indicate that they are larger or smaller than the desired radius, respectively.

$\rho_w(x,y) \propto 1/r_w^2$. The density values are normalized, such that $\max \rho_w = 1$. In practice, we impose an upper bound of 1.3mm on the tube radius. Thus, the density function is bounded from below, with the actual bound varying across the projection surface.

3.2 Computing the disk distribution

Given the density function ρ_w, and a desired number of disks N, we use the method of de Goes et al. [2012] to compute the optimal Voronoi tessellation satisfying the equal capacity constraint, and fit a maximal inscribed disk inside each of the resulting tessellation cells.

Although the above method is able to enforce the capacity constraints exactly, the resulting Voronoi cells are not necessarily hexagonal, and thus some of the maximal inscribed disks may deviate from their intended size. We estimate an initial value for N by computing ρ_d, the amount of density per disk in the dense reference pattern B_0, and dividing the integral of the target density function by ρ_d. We found that this estimate typically results in a large percentage of disks being too small. We thus employ a binary search to find the number of disks N for which the percentage of disks which achieve their intended radius (within 0.05mm) is greatest. Figure 7 shows demonstrates how the percentage of correctly sized disks varies with the total number of disks for the dog example.

While most disks are acceptably close to their desired radius, there are still disks which are too large or too small. Consequently, the amount of light transmitted through the embedded tubes could be smaller or larger than the desired intensity. Larger disks are handled by reducing the embedded tube radius or tilt angle, as needed to match the originally intended ones. Smaller disks are resolved by locally reducing the density in such areas. For each group of such disks, whose Voronoi cells are connected, we remove disks randomly one by one. After each disk removal, we enlarge the group of cells along the one-ring neighborhood to include their immediately neighboring cells, and locally relax the remaining disks in this area via centroidal Voronoi tessellation using Lloyd's method.

3.3 Tube generation

For each disk D centered at (x,y) with radius r_w on S_W, we first subtract the safety margin to obtain a smaller concentric disk D'. In the case where $I^t(x,y) \geq B_0(x,y)$, we generate a tube directed at the center of the light source by intersecting the oblique circular cone defined by D' and the light source center with the lampshade shell, as shown in Figure 8.

If $I^t(x,y) < B_0(x,y)$, a minimal radius tilted tube is generated inside the inner disk D'. In order to avoid structured artifacts, the orientation of the tilted tube inside D' is chosen at random by placing two minimal radius disks D_1 and D_2 near the circumference of D'. D_1 and D_2 are then projected onto the outer and the inner surfaces of the lampshade shell, respectively (see Figure 8). The tilted tube is formed by connecting the two resulting contours. Note that the tilt angle grows with the radius of D'.

At this point, a printable 3D model of the perforated lampshade may be generated and printed by a 3D printer.

3.4 Projected image simulation

The projected image is formed by the light emanating from the light source inside the lamp, passing through the tubes in the lampshade and reaching a receiving surface (wall). Thus, the amount of reflected light at each point on the wall depends on the emission characteristics of the light source, the geometric relationship between the light source and the receiving point, the visibility between them, and the reflectance characteristics of the wall.

The light source that we use in this work is a small and powerful COB (Chips on Board) LED, which is roughly disk-shaped with a diameter of 9mm. In our simulations, we represent this area light source as a collection of n point light sources $\{l_i\}_{i=1}^n$, each emitting a luminous flux of $\Phi_i = \Phi/n$, where Φ is the luminous flux for the entire light source, measured in lumens. We use $n = 76$ in our results.

For a Lambertian light source and a Lambertian reflector, the total illuminance $E_v(p)$ (in lux) reaching a point p on the receiving surface is given by

$$E_v(p) = \Sigma_i \frac{\Phi_i}{\pi r_i^2} \cos(\theta_i) \cos(\theta_p)\, V(p, l_i), \tag{4}$$

where r_i is the distance from p to l_i, and θ_i and θ_p are the angles between the line connecting the two points and the two normals, as shown in Figure 9. $V(p, l_i)$ is the visibility between p and l_i.

In practice, we found that our light source is not Lambertain, and its emitted luminance diminishes as the angle θ_i increases. Furthermore, we found that this behavior is slightly anisotropic, and the luminance diminishes more quickly in the vertical direction. Empirically, we found that the non-Lambertian behavior and the anisotropy are still well modeled by Equation (4), provided that the horizontal and vertical coordinates of the point p (with respect to the center of the image) are each scaled by an appropriate factor.

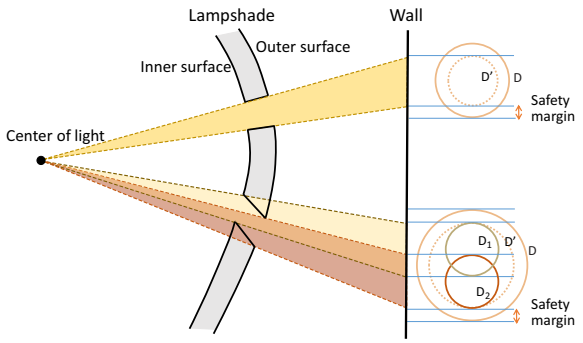

Fig. 8. Illustration of the tube generation.

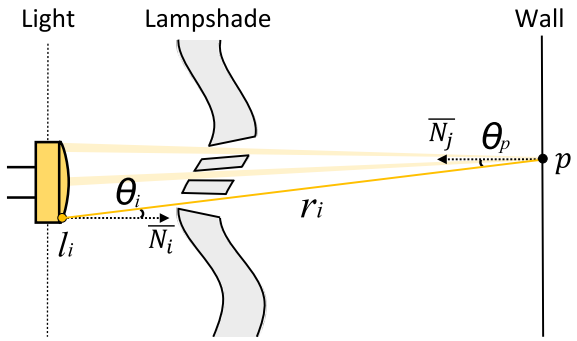

Fig. 9. Illustration of the illuminance computation.

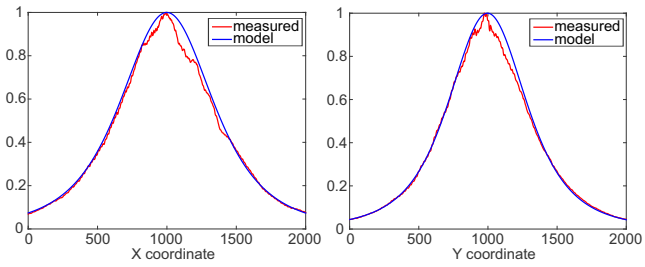

Fig. 10. The measured intensity of the projected image (red) is well predicted by our model (blue). Left: a horizontal slice through the center, right: a vertical one. Note that the horizontal and vertical falloff rates is different.

Specifically, the scaling factors that we use are 1.7 and 1.9 for the horizontal and the vertical coordinate, respectively.

The accuracy of our empirical model is demonstrated in Figure 10. The red curves show a horizontal (left) and a vertical (right) cross section through the projected image shown in Figure 4(c), and the blue curves are the values predicted by our model.

We evaluate the visibility of points on the light source using ray casting (see Figure 9). Note that it is not necessary to intersect rays with the complete geometric model of the perforated lampshade shell. Instead, we quickly determine a small set of tubes that are relevant for a given ray, and check whether the ray passes through one of these tubes, without intersecting the tube's surface. Note that the above approach accounts only for direct illumination, without taking into account reflections, subsurface scattering, etc. In order to assess the magnitude of these effects, we have sprayed the interior of lamps with black paint, and found the differences in the measured intensity very small. The main difference is that this reduces the amount of ambient light in the room, leading to projected images with slightly better contrast.

Accurate light source visibility estimation is particularly important for correctly simulating tilted tubes. Figure 11 shows the simulated and the actual light footprints of a single tube, for a variety of different tilt angles. The sum of the intensity across the area of the footprint is plotted below (normalized such that the maximal sum is 1.0). We can observe that the simulation predicts the illuminance reasonably well, with some deviations between the two curves due to imprecisions in the manufacturing process.

Finally, in order to visualize the simulated results we convert them into grayscale values by scaling the result into a $[0, 1]$ range, and applying a gamma correction (we use gamma 2.2 in all our results).

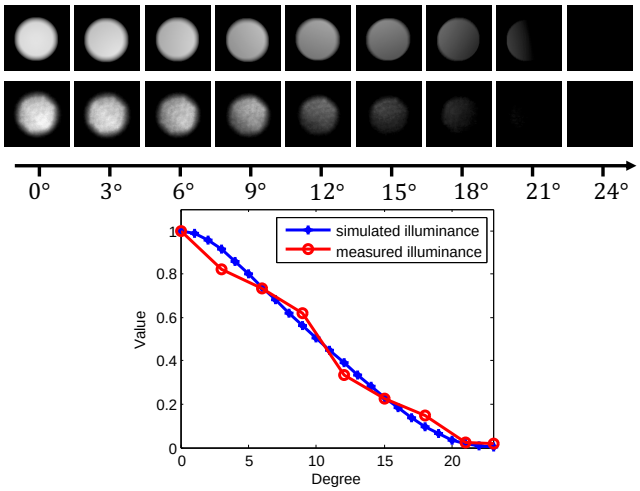

Fig. 11. One tube (diameter 1.2mm) with varying tilting angles. Top row: simulated footprints; Second row: the physical footprints; Bottom: plot of the normalized sum of simulated illuminance values in the footprint and the sum of grayscales across an actual footprint photo.

4. RESULTS

We have implemented the method described in the previous section, and have used the resulting system to design a number of lampshades for projecting a variety of target images. The results may be seen in Figures 1 and 16. The lampshades in all of these results have spherical geometry, with a diameter of 22cm. To save material we typically perforate each spherical lamp with two images on two opposing sides. The light source we use in our results is a Cree® XLamp® CXA1507 LED with 3000K color temperature, which is roughly disk-shaped with a radius of 9mm. The perforated lamp projects its image onto a planar rear projection screen that offers high light transmission, located at a distance of 40cm from the light source. The physical size of the projected image is 100 × 100 cm. Figure 12 shows a photograph of our setup. The photographs of the projected images shown in this paper are shot from the opposite side of the screen for eliminating distortion, using a Canon EOS 5D Mark II camera, with exposure 1/80 sec, f/4.0, and ISO 400 settings, in a room without any additional light sources.

Fig. 12. A photograph of our setup.

Table I. Statistics for our results

Model	Initial #disks	Final #disks	% correct size
Marilyn	7323	5789	79.47
CircularRamp	7914	6321	87.73
Hepburn	7617	5945	82.50
Kelly	7563	5914	80.74
Marlon	7167	5607	81.34
Toucan	6926	5352	75.76
Lion	7243	5650	74.81
Dog	7891	6055	82.12
Kelly (bumpy)	8912	7120	65.17

Initial disks is our initial estimate of the required number of disks, which serves as a starting point for a binary search, after which we achieve the final number of disks. The percentage of correctly sized disks among these are reported in the right column.

The resulting lampshade designs have been 3D-printed to allow a qualitative evaluation of the projected imagery. The lampshades were printed on a powder-based binder-jet printer (ProJet® 660Pro). The thickness of the lampshade shells was set to 3mm, which, in our experience, is the lower bound to ensure the structural integrity. We found that it is hard to blow the powder away from the tubes when $r_{\min}$ is smaller than 0.6mm, and that adjacent tubes are prone to collapse when $d_{\min}$ is smaller than 0.5mm. Therefore, we set $r_{\min} = 0.6$mm and $d_{\min} = 0.5$mm. The manufacturing time for such a spherical lampshade includes 16.5 hours of printing and 1 hour of drying. The inside surfaces of the lamps were sprayed with black paint, which somewhat improves the contrast of the projected images.

Our method is also applicable to non-spherical lampshades, as demonstrated in Figure 17, where the lampshade is a bumpy surface (bounded by a sphere of diameter 22cm). This is made possible by the fact that the capacity-constrained Voronoi tessellation takes place on the planar receiving surface, rather than on the more general lampshade surface. The tubes are generated by back-projecting disks from the receiving surface through the lampshade shell, which can thus be of a general, non-spherical shape. It should be noted, however, that our current implementation assumes that the lampshade shape is free from self-occlusions.

Table I reports several statistics for each of these results. The computation time of our method consists of two main parts: computing the disk distribution and generating the 3D model of the lampshade perforated with tubes. For the lamps shown in the paper it takes less than 5 minutes to run the method of de Goes et al. [2012] up to 10 times, while searching for the optimal number of disks. The resulting lamps have around 6000 tubes, and the 3D model generation takes roughly 1 minute. For verification, we typically compute a simulated result before sending a lamp to be printed, which takes roughly 1 minute. All of the above times are measured on an Intel® Core™ i5 CPU 3.3GHz with 8GB RAM.

Figure 13 shows two qualitative comparisons between the target image, the simulated result, and the physical projected image. This is done on two examples: a concentric circles pattern featuring four regions of a constant grayscale value, and a more complex natural image. Below each image, we show the corresponding histogram, in order to better examine the differences between the corresponding tone distributions.

In the concentric ramp case (13a), it obviously infeasible to reproduce the ideal four-way distribution of grayscale values with our process. Our design process approximates this distribution with four wider peaks (the black peak coincides with the left edge of the histogram). A photograph of the actual projection features four

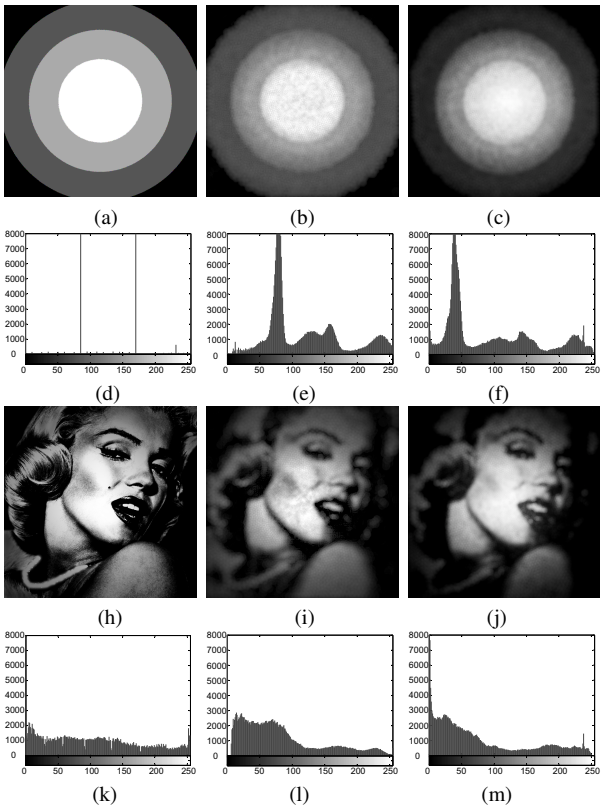

Fig. 13. Two comparisons of the target image (a,h), our simulation result (b,i), and a real photo of the projected image (c,j). The corresponding histogram is below each image. The circular ramp (a) consists of three gray values: 85, 170, 255.

similar peaks, although there are slight differences in their positions and spread. The target image of Marilyn Monroe (13h) has a nearly uniform histogram. This histogram is reproduced reasonably well by the our process, however, both the simulated and the actual projected image have more mass in the lower part of the grayscale range.

These two examples illustrate well some of limitations of this new medium. Achieving bright tones in the middle of the image is relatively easy (since the distance between the receiver and the light source is minimal there, and both cosines in Eq. (4) are near 1.0). However, reproducing dark tones in the central region is challenging because of the necessity to use large tilt angles. As explained earlier, extensive use of tilting requires reducing the tube density, thereby reducing spatial resolution. Conversely, achieving light tones at the periphery of the image is difficult, because of the falloff in the illuminance, and increasingly wider tubes are needed to produce a light tone in the periphery. Thus, there's a tradeoff between our ability to reproduce the full tonal range, and our desire for better spatial resolution. In particular, by imposing an upper bound (1.3mm) on the embedding disk radii across the lamp surface, we are not able to match the lighter tones of the target image beyond a certain radius.

The above can be observed in a more quantitative manner in Figure 14, which shows the response of our method to several radial cosine waves with the frequency increasing from top to bottom row. Next to each input pattern we show the simulated result produced by our method and plot the amplitudes of the input wave (in red)

and the result (blue), as a function of the radial distance. It may be observed that for all frequencies our method fails to reproduce the waveform beyond a certain radial distance, with the distance slightly decreasing as the frequency increases. It may also be seen that in the areas where the waveform is reproduced, the contrast response also decreases with frequency: while in the top row the ratio between the reproduced amplitude to the input one is roughly 0.89, this ratio drops to 0.40 in the bottom row.

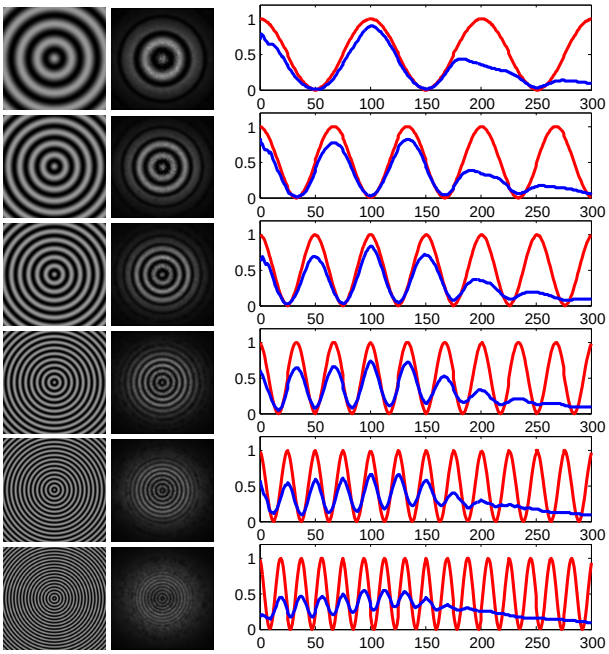

Fig. 14. Frequency and contrast response. Given the radial cosine waves as the input patterns (left), the results of our method (middle) are influenced by both frequency and radial distance from the center, as shown in the amplitude plot (right).

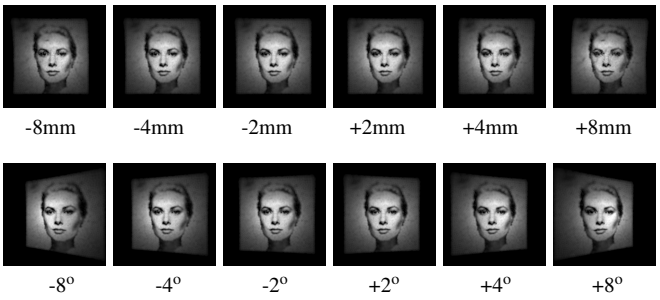

Fig. 15. Simulated projected images with the position of the lampshade shifted horizontally (top row) and lampshade rotated (bottom row) for the Kelly model shown in Figure 17.

In practical usage, the position of the lampshade or the inside light source, or the relative orientation between the two may differ slightly from the ideal conditions assumed in the design process. Figure 15 demonstrates that the projected image is still acceptable, so long as these deviations are not too large.

5. CONCLUDING REMARKS

We have presented a technique to design printable perforated lampshades. The main challenge was to control the light emanating through the lampshade under the unique printability constraints and engineering setting. The technique that we developed extends classical halftoning to a novel domain, where the "dots" are not printed, or physically tangible, but projected, effectively *halftoning with light*.

Our current implementation assumes that the receiving surface is planar, and that there are no self occlusions due to the shape of the lampshade. Thus, the distribution and the geometry of the perforating tubes can be determined by intersecting cones formed by disks on the receiving surface and the center of the light source with the lampshade shell. However, our framework is able to support more general lampshade shapes and receiving surfaces. Specifically, the generation of tubes needs to be adjusted to ensure that light reaches the receiving surface even in areas where self-occlusions exist, and the transformation between areas on the receiving surface and tube cross sections must be computed in a more general fashion.

The 3D printing technology is emerging and growing quickly, and many new intriguing applications are introduced. We believe that the combination of light and 3D printing has much more to offer. One direction which we now consider is not to use the light directly, but indirectly, by using the printed surface as reflector. The printed surface can be customized to the given environment, to its geometry and photometric properties, so as to optimize the distribution of light by a custom-made reflector.

Another direction for future work is to go beyond the printability constraints, by assembling a large scale lampshade from surface pieces printed separately. This will allow using bigger and stronger light sources, and possibly compound arrays of light sources. A larger surface area will allow increasing the relative resolution of printed holes, and project an image on larger and more distant receiving surfaces. We are also considering focusing on the design of aesthetic imagery on the lampshade itself. Halftoned and stippled images have their own aesthetic virtue. In Figure 18, we show our lamp in a bedroom. Light effects combined with halftoning can lead to creative artistic media.

ACKNOWLEDGMENTS

This work is partially supported by China National 973 Program (2015CB352501), Natural Science Foundation of China (61232011, 61572291, 61332015), NSFC-ISF (61561146397), and the Israel Science Foundation (ISF).

REFERENCES

ALEXA, M. AND MATUSIK, W. 2012. Irregular pit placement for dithering images by self-occlusion. *Computers & Graphics 36,* 6, 635 – 641. 2011 Joint Symposium on Computational Aesthetics (CAe), Non-Photorealistic Animation and Rendering (NPAR), and Sketch-Based Interfaces and Modeling (SBIM).

BALZER, M., SCHLÖMER, T., AND DEUSSEN, O. 2009. Capacity-constrained point distributions: A variant of Lloyd's method. *ACM Trans. Graph. 28,* 3 (July), 86:1–86:8.

BARAN, I., KELLER, P., BRADLEY, D., COROS, S., JAROSZ, W., NOWROUZEZAHRAI, D., AND GROSS, M. 2012. Manufacturing layered attenuators for multiple prescribed shadow images. *Comp. Graph. Forum 31,* 2 (May), 603–610.

CHANG, J., ALAIN, B., AND OSTROMOUKHOV, V. 2009. Structure-aware error diffusion. *ACM Trans. Graph. 28,* 5 (Dec.), 162:1–162:8.

Fig. 16. Additional results on spherical lampshades. Left: our 3D-printed lampshade, with light projecting through the surface. Right: the photo of projected image. The lamps are 22cm in diameter.

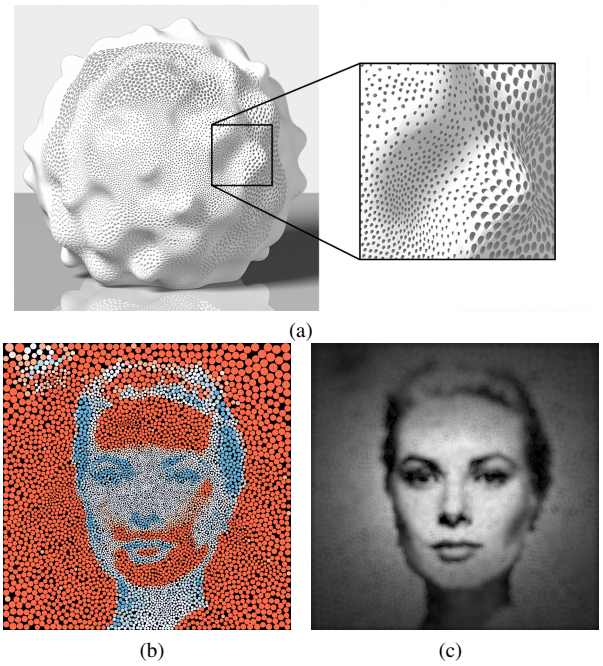

Fig. 17. A non-spherical lampshade (a), the maximal inscribed disks inside each cell of the CCVT (b), shown using the same color coding as in Figure 2, and the resulting simulated projected image (c).

Fig. 18. One of our manufactured lamps in a natural setting.

DE GOES, F., BREEDEN, K., OSTROMOUKHOV, V., AND DESBRUN, M. 2012. Blue noise through optimal transport. *ACM Trans. Graph. 31,* 6 (Nov.), 171:1–171:11.

ESCHBACH, R. AND KNOX, K. T. 1991. Error-diffusion algorithm with edge enhancement. *J. Opt. Soc. Am. A 8,* 12 (Dec), 1844–1850.

FINCKH, M., DAMMERTZ, H., AND LENSCH, H. 2010. Geometry construction from caustic images. In *Computer Vision - ECCV 2010*, K. Daniilidis, P. Maragos, and N. Paragios, Eds. Lecture Notes in Computer Science, vol. 6315. Springer Berlin Heidelberg, 464–477.

KIPPHAN, H. 2001. *Handbook of Print Media: Technologies and Production Methods*. Springer-Verlag.

KISER, T., EIGENSATZ, M., NGUYEN, M., BOMPAS, P., AND PAULY, M. 2013. Architectural caustics – controlling light with geometry. In *Advances in Architectural Geometry 2012*, L. Hesselgren, S. Sharma, J. Wallner, N. Baldassini, P. Bompas, and J. Raynaud, Eds. Springer Vienna, 91–106.

LAN, Y., DONG, Y., PELLACINI, F., AND TONG, X. 2013. Bi-scale appearance fabrication. *ACM Trans. Graph. 32,* 4 (July), 145:1–145:12.

LAU, D. L. AND ARCE, G. R. 2008. *Modern Digital Halftoning, Second Edition (Signal Processing and Communications)*, 2 ed. CRC Press.

LEVIN, A., GLASNER, D., XIONG, Y., DURAND, F., FREEMAN, W., MATUSIK, W., AND ZICKLER, T. 2013. Fabricating BRDFs at high spatial

resolution using wave optics. *ACM Trans. Graph. 32,* 4 (July), 144:1–144:14.

LI, H., WEI, L.-Y., SANDER, P. V., AND FU, C.-W. 2010. Anisotropic blue noise sampling. *ACM Trans. Graph. 29,* 6 (Dec.), 167:1–167:12.

LOU, Q. AND STUCKI, P. 1998. Fundamentals of 3D halftoning. In *Electronic Publishing, Artistic Imaging, and Digital Typography, 7th International Conference on Electronic Publishing, EP '98, Held Jointly with the 4th International Conference on Raster Imaging and Digital Typography, RIDT '98, St. Malo, France, March 30 - April 3, 1998, Proceedings*. 224–239.

MALZBENDER, T., SAMADANI, R., SCHER, S., CRUME, A., DUNN, D., AND DAVIS, J. 2012. Printing reflectance functions. *ACM Trans. Graph. 31,* 3 (June), 20:1–20:11.

MITRA, N. J. AND PAULY, M. 2009. Shadow art. *ACM Trans. Graph. 28,* 5 (Dec.), 156:1–156:7.

MITSA, T. AND PARKER, K. J. 1992. Digital halftoning technique using a blue-noise mask. *J. Opt. Soc. Am. A 9,* 11 (Nov), 1920–1929.

PANG, W.-M., QU, Y., WONG, T.-T., COHEN-OR, D., AND HENG, P.-A. 2008. Structure-aware halftoning. *ACM Trans. Graph. 27,* 3 (Aug.), 89:1–89:8.

PAPAS, M., HOUIT, T., NOWROUZEZAHRAI, D., GROSS, M., AND JAROSZ, W. 2012. The magic lens: Refractive steganography. *ACM Trans. Graph. 31,* 6 (Nov.), 186:1–186:10.

PAPAS, M., JAROSZ, W., JAKOB, W., RUSINKIEWICZ, S., MATUSIK, W., AND WEYRICH, T. 2011. Goal-based caustics. *Comp. Graph. Forum 30,* 2 (June), 503–511.

PEREIRA, T., RUSINKIEWICZ, S., AND MATUSIK, W. 2014. Computational light routing: 3D printed optical fibers for sensing and display. *ACM Trans. Graph. 33,* 3 (June), 24:1–24:13.

SCHWARTZBURG, Y., TESTUZ, R., TAGLIASACCHI, A., AND PAULY, M. 2014. High-contrast computational caustic design. *ACM Trans. Graph. 33,* 4 (July), 74:1–74:11.

STUCKI, P. 1997. 3D halftoning. In Proc. SPIE 2949, Imaging Sciences and Display Technologies. *Proc. SPIE 2949*, 314–317.

ULICHNEY, R. 1987. *Digital Halftoning*. MIT Press, Cambridge, MA, USA.

ULICHNEY, R. 1988. Dithering with blue noise. *Proceedings of the IEEE 76,* 1 (Jan), 56–79.

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,* 4 (July), 95:1–95:12.

WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Trans. Graph. 28,* 3 (July), 32:1–32:6.

WILLIS, K., BROCKMEYER, E., HUDSON, S., AND POUPYREV, I. 2012. Printed optics: 3D printing of embedded optical elements for interactive devices. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*. UIST '12. ACM, New York, NY, USA, 589–598.

YUE, Y., IWASAKI, K., CHEN, B.-Y., DOBASHI, Y., AND NISHITA, T. 2014. Poisson-based continuous surface generation for goal-based caustics. *ACM Trans. Graph. 33,* 3 (June), 31:1–31:7.

ZHOU, C. AND CHEN, Y. 2009. Three-dimensional digital halftoning for layered manufacturing based on droplets. *Transactions of North American Manufacturing Research Institute of SME 37*, 175–182.