Computers & Graphics 37 (2013) 974–982

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Section on Advanced Displays

Parallax Walls: Light fields from occlusion on height fields

Xavier Snelgrove*, Thiago Pereira, Wojciech Matusik, Marc Alexa

223 Queen St. E, Toronto, ON, Canada M5A 1S2

ARTICLE INFO

Article history:
Received 16 April 2013
Received in revised form
29 June 2013
Accepted 11 July 2013
Available online 22 August 2013

Keywords:
Light fields
Reflectance fields
Parallax
Passive displays

ABSTRACT

We introduce a simple, inexpensively manufacturable, height field surface geometry that, when appropriately painted, can produce diffuse reflections under ambient lighting approximating a target light field. We demonstrate a light field basis analysis of these surfaces which allows us to formulate the problem as a gamut-mapping, and propose a perceptually motivated metric for bringing complex light fields into gamut while preserving their structure. We show resultant surfaces displaying physical light fields, animations, HDR exposure stacks, and scene relighting.

1. Introduction

Progress on active displays allows us to show dynamic content that is sensitive to view point (autostereoscopic) [1] or even the lighting environment [2]. Yet the world is primarily passive: book covers, paintings, billboards and even buildings are examples of passive surfaces that could be used as displays (Fig. 2). While these passive displays do not change their content as a function of time, they can still show complex content that varies with view point or lighting.

Passive view-sensitive and light-sensitive displays fail to scale to large setups due to geometric complexity. For instance, using lenticular arrays for building a billboard-sized display would either need very high resolution printing or prohibitively large lenslets. Parallax barriers are also not suited for passive displays due to the need of using only incoming environment light.

In this work, we focus on the concept of occlusion for conveying several different views. In particular, we depart from the notion that barriers are necessarily parallel to the display plane. Instead, we focus on *walls* that are *orthogonal* to the surface. This idea has several interesting consequences:

- The barriers are much easier to fabricate since the resulting surface is a height field. We also note that a variety of “grills” are readily available from other industries.
- Nearly all light directed towards the surface is reflected. Consequently, the resulting images are much brighter which is necessary for passive displays.
- The type of light field that can be displayed has certain restrictions, which are different from those imposed by classical parallax barriers. Clearly, as the view direction departs from the centre, one strictly sees a subset of the surface. This means off-centre views are always darker than central views.

A major contribution of this work is the *analysis* of displays based on occlusion in terms of basis functions (*i.e.* matrices) for different barrier directions (see Section 3). This analysis clearly shows advantages and disadvantages of different display types: parallax barriers simply reduce the brightness, orthogonal barriers shear this gamut, so that images can be brighter, at the cost of greater correlation.

Our analysis also readily provides a discretization of view directions. Based on this discretization, we can optimize static displays for given light fields using non-negative optimization techniques (Section 4). Because of the limited gamut, we explore different techniques to exploit the human visual system and deficiencies for better mappings into the space provided by the displays.

Apart from stereoscopy, the orthogonal barriers naturally have other applications. By exchanging a moving viewer by a moving light source, we show how parallax walls can also be effectively used to show relightable content. We show many relighting applications where the images are naturally darker for large light incident angles which means our gamut limitations introduce very mild distortions. Because of the high brightness, we can also use this type of display for high dynamic range exposure stacks. More results and applications are shown in Section 5.

This new type of occlusion for autostereoscopy-type applications has great potential, as it provides cheaper production, more

* Tel.: +1 416 876 9427; fax: +1 647 556 2109.
E-mail address: xavier@cs.toronto.edu (X. Snelgrove).

0097-8493/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cag.2013.07.002

brightness, at the cost of darker images at grazing angles, and correlation requirements on the views. Our proposed design does not have the same shortcomings of other approaches for extremely large-scale applications as shown in Fig. 2.

2. Related work

Microgeometry for custom appearance: Generating surface geometry to match desired appearance properties has been investigated recently. For instance Weyrich et al. [3] investigate the effect of specular microfacets on macro-scale specularities. Our approach relies on the effects of self-occlusion or self-shadowing, which have also been exploited to display a single custom image [4], or multiple images based on lighting direction [5,6]. All these techniques require milling, a technique which would not scale to large objects.

Interaction of surface geometry and material: The relationship of surface geometry and texture to appearance has been investigated extensively for areas such as BRDF modelling [7], and stereographic image-pair display [8]. Recent work uses periodic surface geometry with specular paint to display relightable objects [9]. Our combination of diffuse surface texture with simple geometry for light field display is novel.

Attenuators for view dependent display: The use of attenuators to create view dependent display has been investigated for light field display [1,10], or lighting-dependent shadows [11]. This work often falls under the heading of "computational displays" [12]. Our design, however, uses more easily manufacturable *binary* attenuators, which have also been used previously for lighting-dependent shadows [13].

Gamut mapping: A major contribution of this paper is an achievable-gamut based analysis of surface appearance. The problem of taking out-of-gamut content and mapping it into gamut in a perceptually appropriate way has been investigated for other problem domains, for instance mapping HDR images to LDR displays [14], or mapping general 3d content to a flatter representation for bas-relief design [15].

To the best of our knowledge, our gamut-based analysis is novel in the context of light field display, although analogous to the depth-dependent low-pass filtering of content for light field displays to prevent aliasing [1], or the warping of target shadows to create binary volumes with those shadows [13].

3. From height fields to light fields

The surface we propose is a simple height field, shown in Fig. 1. We consider painting the bottom face of this surface with diffuse inks. Each point on the surface is visible from a range of angles. In this section, we analyse the images seen from different viewing positions. Keep in mind that this same analysis holds for the case of multiple lighting positions and a fixed viewer.

3.1. Continuous light field basis

We analyse this surface by the *light field* it produces. The light field is a 4D function measuring the radiance along rays in free space [16,17]. A convenient parameterization of the light field is the *two-plane parameterization*, where the intersection points of a ray with two parallel planes uniquely determine that ray. Different views of the surface correspond to 2d slices of this 4d space.

We consider only a 3d slice of this 4d light field, corresponding to the horizontal parallax in the scene, and to simplify the representation we place one of the planes on the back face of the surface, and one at infinity. Our light field then is $L(x, y, u)$ where x and y are the intersection points of the ray with the back plane, and u corresponds to a ray direction (homogeneous coordinates can be used to deal with the u-plane being at infinity as in [16]).

Each drop of diffuse paint creates a basis line in the light field, as shown in Fig. 3. Different views of the scene correspond to different slices of this light field, and orthographic views "from infinity" correspond to slices of constant u. In a standard parallax barrier design with infinitesimal opening size these basis lines do

Fig. 2. Rendering of our design as it might appear on the side of a building. Here as the sun travels across the sky, the lighting of the face changes realistically, despite the display being passive.

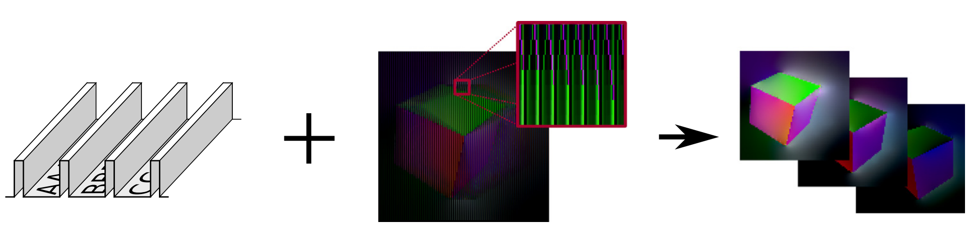

Fig. 1. Overview of our system. We produce a surface with a cross-section as in the left-most image, and print a texture onto the flat bottom-portions of the surface. The texture shown in the centre image produces the light field at right when painted on the surface.

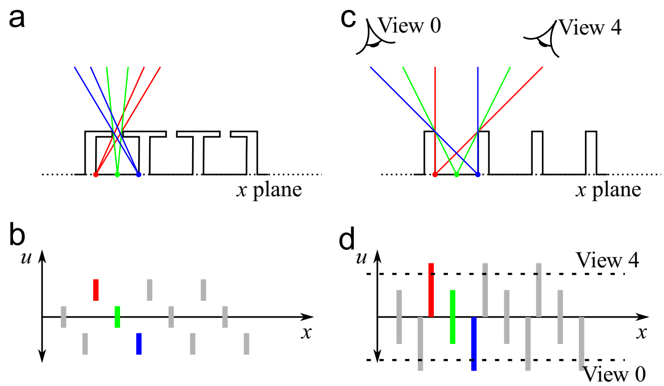

Fig. 3. Comparison of light fields generated by paint spots at various points for both our design, and a standard parallax barrier. The upper images show a finite-aperture parallax barrier and our design respectively, as well as the rays emitted by three points on the surface painted red, green, and blue. The bottom shows these same rays in a 2d constant-y slice of the light field. Each bundle of rays in the upper image maps to a line in the lower. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this article.)

not overlap in u: the views are independent. This results in masking out most of the light leaving the surface. In our design much less light is masked out, but the views are more correlated. The lines from Fig. 3 are the basis with which we work to create light fields of our choosing. The coefficients are the surface colours, and so must be non-negative.

Note that the vertical walls in the parallax barrier pictured are not standard of parallax barrier designs, we include them to make the comparison more readily apparent. In a standard parallax barrier there is view repetition, where the light field shown in Fig. 3(b) has infinite repeated copies shifted in x and u. For a limited field of view, however, the formulations are identical.

3.2. Discretizing the light field basis

We perform the bulk of the analysis in a discrete domain. Our light field is defined over (x, y, u), which we discretize to a resolution of $X \times Y \times U$. The surface is naturally discretized in x by its walls, which divide the surface into cells. For simplicity we model each cell as having a single colour at a given y-coordinate for a given view. This colour is the average of the surface colours visible in that view. We define a *hogel* as the vector of brightness in all views of a point at a given (x, y) coordinate. This term is borrowed from an analogous quantity in holography [18].

We can set Y arbitrarily up to the printer resolution, although it is typically set to the value ensuring that our hogels are square in x and y. Finally we discretize U. In Fig. 3(c) we have marked three distinct regions per hogel in blue, green and red, this corresponds to a u-resolution of $U = 3$. Let the colours of these points be a_0, a_1, and a_2 (note that we have labelled the colours from right to left). Consider the left-most view (v_0) from which the bottom of the cell is still visible. This view can see only a_0 (the blue region). In other words

$$v_0 = \left[\tfrac{1}{3}\ 0\ 0\right][a_0\ a_1\ a_2]^{\mathrm{T}} \tag{1}$$

The next view is the sum of the blue and green regions a_0 and a_1. Continuing this logic gives a set of 5 discrete views:

$$\begin{bmatrix} v_0 \\ v_1 \\ v_2 \\ v_3 \\ v_4 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ a_2 \end{bmatrix} \tag{2}$$

subject to the constraints:

$$0 \le a_i \le 1 \tag{3}$$

This normalizes the brightest possible view to a value of 1. We can write this more compactly as follows:

$$\mathbf{v} = B\mathbf{a} \tag{4}$$

Note that although in this example $U = 3$ the basis matrix, B, has 5 rows. However it is of rank 3, so the 5 views are not independent. As such we may as well reduce it to a square matrix by selecting the three views we care about. Natural choices would be three views around the centre (so the middle three rows of B), or three views from the centre off to one side (the last three rows of B).

$$B_c = \frac{1}{3} \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix} \tag{5}$$

$$B_s = \frac{1}{3} \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \tag{6}$$

For comparison, the basis matrix for a standard parallax barrier, as shown in Fig. 3(a), is

$$B_{par} = \frac{1}{3} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{7}$$

where the $\frac{1}{3}$ scaling factor in the above is due to this example being $U = 3$. Our display of course supports values of U up to the printer resolution; in general the scaling factor will be $1/U$.

The angular extent of each view is determined by the resolution X, as well as by the height of the occluding walls. If the walls are very high, the views are compressed into a very small angular extent, and conversely if the walls are very short, the views cover a wide range of angles. This is apparent from Fig. 3(c). Wall height also determines the amount of ambient irradiance onto the surface, so as wall height increases, brightness decreases. Our image formation matrices are normalized by the brightness of the central view, and each row corresponds to a view, not to an angle, so changing wall height does not affect those equations at all.

3.3. Achievable gamuts

The matrices B_s and B_c are full-rank, however the constraints of Eq. (3) mean that not all hogels are achievable. Comparing the achievable gamut of hogels between various approaches illuminates some differences between the designs. Fig. 4 compares the standard parallax barrier, our design around the centre as in Eq. (5) and our design off to one side as in Eq. (6). We select $U = 3$ as it allows the gamuts to be visualized in 3-space.

There are a few quantities of note in these achievable gamuts. Although all three have the same volume, which is the determinant of the B matrix, however the brightest achievable points vary. The parallax barrier is the darkest solution, with a maximum brightness of $(1/U, 1/U, 1/U)$. The next brightest is the offset solution B_s, with a maximum brightness of $(1, 2/U, 1/U)$. Finally the centred solution B_c is brightest at $(2/U, 1, 2/U)$.

This additional brightness does come at a cost: view correlation. This is the *brightness-correlation tradeoff*. A sheet of regular paper is one extreme of this tradeoff, where the views are as bright as possible, but completely correlated. The traditional parallax barrier is the other extreme, the darkest but least correlated solution. Note that in the traditional parallax design (the first column of Fig. 4) the opening angle of the achievable gamut is such that the three views are completely uncorrelated. As we

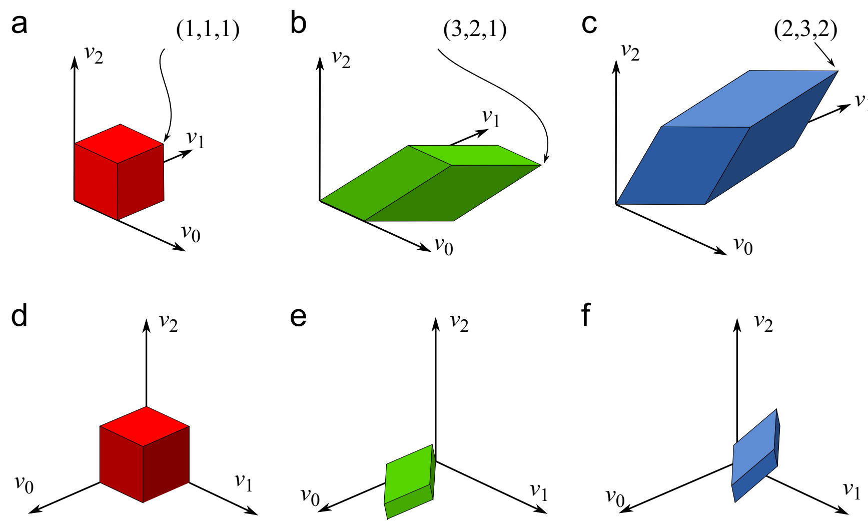

Fig. 4. Achievable gamuts for different bases with $U=3$. The red, leftmost, column shows the parallax barrier gamut B_{par}. The green, middle, column shows the gamut for B_s and the blue, rightmost, column shows the gamut for B_c. The top and bottom rows are rotations of the same gamuts, and the labelled values have been scaled by U. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this article.)

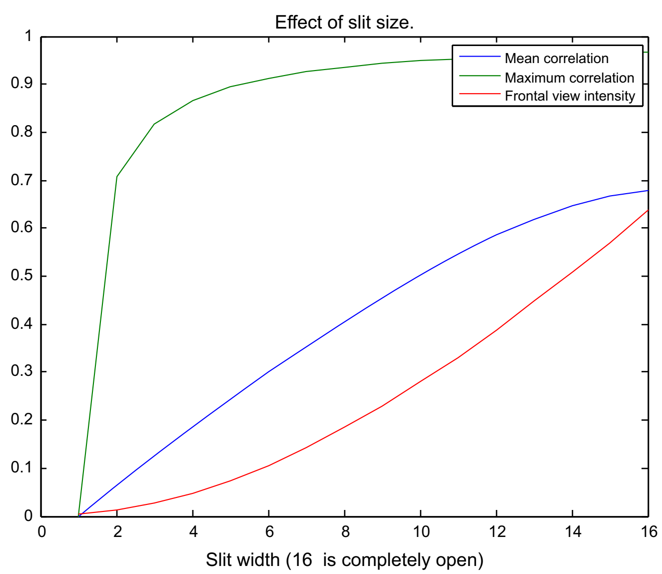

Fig. 5. Impact of slit size, slit size measured in pixels, with 16 pixels between walls. Frontal view intensity (red line) increases quadratically with slit size since not only more pixels are visible but also more ambient light is allowed in. Maximum correlation (green) increases very rapidly since neighbouring views differ by only one pixel. On the contrary, mean correlation (blue) increases very slowly. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this article.)

stretch the brightest point further from the origin, this opening angle becomes narrower: the views become more correlated. Our design falls along this curve, acknowledging that natural light fields are highly correlated [1], so the full power of the parallax barrier is not required.

In Fig. 5 we show the impact of varying the slit size (as in Fig. 3) on both view correlation and maximum intensity. In this simulation, a slit of size one pixel corresponds to a parallax barrier and a fully open slit of size 16 corresponds to our proposed parallax walls. We simulated with a height of 13 pixels as used in our physical prototypes.

Intensity has been normalized where one is the value of intensity of perfectly diffuse paper under ambient illumination. When decreasing the slit size, intensity decreases for two reasons: fewer pixels are visible from each view and more incoming ambient light is blocked. Notice how even parallax barriers with a big slit size are at least five times darker than diffuse paper, while parallax walls can achieve 60% of paper intensity. We remind the reader that this comes at the cost of view correlation.

For the correlation analysis, we assume random pixel values between 0 and 1. As such, each view v_i becomes a random variable. We calculate all the correlations between two different views and plot the mean and maximum values. We can see how the maximum correlation increases very rapidly since neighbouring views differ by only one pixel. On the contrary, mean correlation increases very slowly which better explains how two distant views can display different content.

Note on ambient occlusion: In this analysis we treat each point on the surface as being equally brightly illuminated. In the case of ambient lighting this will not likely be the case, due to ambient occlusion on the points adjacent to the walls. For a known illumination condition, then, this matrix B would be slightly re-weighted.

3.4. Comparison to other techniques

3.4.1. Parallax barriers

Although we compare favourably to parallax barriers from a brightness perspective as discussed above, there is one big advantage to our approach that this analysis does not cover. A parallax-barrier surface is darker not only because most of the surface radiance is masked (by design) but also because most of the surface *irradiance* (incident light power per unit area) is masked. This means that parallax-barriers are well-suited to backlit displays, but not to diffuse surfaces lit by ambient light. Our approach, on the other hand, allows for passive surfaces.

3.4.2. Lenticular arrays

Lenticular arrays in which small lenses, rather than occlusion barriers, are used to direct light in the desired direction do not suffer from the brightness-correlation tradeoff of the approaches discussed above. A lenticular display, in theory, can show uncorrelated views at full brightness. In practice there is still cross-talk between lenses because the PSF of the lenses is not a point, as the lenses are typically of low quality. Our approach, however, allows for simpler and cheaper manufacturing as no refractive elements are required. In particular, large setups can be cheaply realized with occlusion barriers.

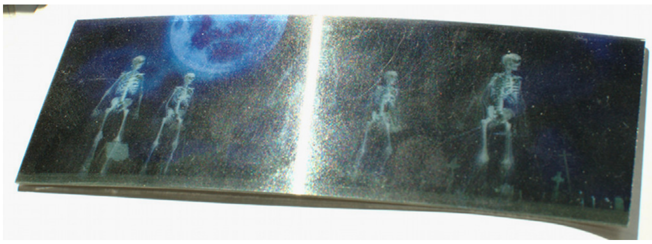

Fig. 6. A lenticular surface on a sunny day. The circular cross-section of the lenses means that some point on every lens reflects the sun directly into the eye.

For large-scale applications such as billboards or architectural features as shown in Fig. 2 lenticular arrays are not well suited. The lenticular lenses are volumetric entities, so their volume increases with the square of hogel width, leading to prohibitively large lenses, unsuitable to being built into the side of a building.

Further lenticular arrays have a field-of-view limited by small-angle assumptions in the lens designs. Our approach makes no such assumptions, so can support much wider fields of view. Lenticular displays are highly specular, due to their smooth optical surfaces. This can cause the display to be unusable under many common lighting conditions with strong point sources as shown in Fig. 6.

Finally, our approach further supports passively relightable scenes, as discussed in Section 5, an application to which lenticular surfaces are not well suited.

4. Optimization

In general we set up the optimization as follows. We have a target light field to be displayed T, where $\mathbf{t}(x,y)$ refers to the vector of views, *i.e.* hogel, at point (x,y). The notation $t_u(x,y)$ refers to the uth view of that hogel.

A given surface is painted with colours $\mathbf{a}(x,y)$ at each surface point (x,y). This generates a hogel $B\mathbf{a}(x,y)$ at that point. We denote its uth view $[B\mathbf{a}(x,y)]_u$.

4.1. Preserving intensities

Perhaps the simplest way to determine a texture to print on our surface is by finding the in-gamut light field that is closest in a least-squares hogel-distance sense to the target light field, that is: minimize

$$\arg\min_{\mathbf{a}} E(\mathbf{a}) = \arg\min_{\mathbf{a}} \sum_{x,y,u} (t_u(x,y) - [B\mathbf{a}(x,y)]_u)^2 \quad (8)$$

subject to $0 \le a_i(x,y) \le 1$. Here $\mathbf{a}(x,y)$ is the vector of the colours painted on the surface at $(x,\ y)$ generating a hogel with views $B\mathbf{a}(x,y)$.

This approach works well if the target is expected to be mostly within gamut. For instance, our B_s basis in Eq. (6) has the property that views get monotonically darker as the viewer move further right. Exposure stacks for HDR have an analogous property: the image gets monotonically darker with shorter and shorter exposures. So displaying an HDR image on the surface such that different viewing angles show different exposures as in [19] is very well-suited to our design, and the hogel-distance optimization approach works well (see Fig. 12c). Some relighting examples are also mostly in gamut.

Each hogel's vector of coefficients, $\mathbf{a}(x,y)$, can be independently solved to minimize its term in the objective function, so this approach is very fast.

4.2. Preserving edges

For targets where that cannot be assumed to have content mostly in gamut, the hogel distance optimization outlined above is not the best approach. Consider a simple 2d light field of a cyan-coloured object some distance away from the camera, as shown in Fig. 7(a). The strong sloped lines are typical of light fields, and their slope determines the depth of the object. In this example the right-hand edge is out of gamut, as increasing u does not monotonically decrease the brightness.

The closest in-gamut solution, in a least-square sense, is shown in Fig. 7(b). Unfortunately the out of gamut edge has been blurred out and none of its oriented structure remains. This is a problem for displaying 3d geometry. Fig. 7(c) shows a more desirable in-gamut light field, where although colour was not preserved, edge structure was, and so by extension geometry. This implies that an optimization which preserves gradients is desirable.

For edge preservation, we could minimize the difference of gradient vectors but this solution can lead to strong gradient compression. We chose instead to minimize gradient magnitudes which let us achieve stronger edges (Fig. 7(c)) with the drawback that they might be flipped. A general form for such an optimization is

$$\arg\min_{\mathbf{a}} E(\mathbf{a}) = \arg\min_{\mathbf{a}} \sum_{x,y,u} |(|\nabla t_u(x,y)| - |\nabla [B\mathbf{a}(x,y)]_u|)|^p \quad (9)$$

The value of p sets which norm to minimize over. For light fields we have had good results with both the $p=1$ and $p=2$ norms, depending on scene content (see Section 6).

Note that this formulation does not explicitly tie the different *views* together. As long as the gradients in x and y are accurate, the views could look completely different. However B implicitly requires that views be correlated, so in practice this is not a problem.

Choice of gradient: Although in (9) we show the spatial gradient in the image as ∇, in general this can be a perceptually motivated distance function between neighbouring hogels. This also allows us to use the same formulation for colour images: in our case we use the CIELAB metric for an approximation of perceptual colour distance, and then let $\nabla_x t_u(x,y) = \| l_u(x+1,y) - l_u(x,y) \|$ where $l_u(x,y)$ is the CIELAB representation of the colour $t_u(x,y)$. The formulation for ∇_y is similar.

Fig. 8 shows the difference between using the gradient-magnitude preserving objective and a simple hogel-distance metric. Clearly for out-of-gamut light fields such as this, our perceptually motivated objective is superior.

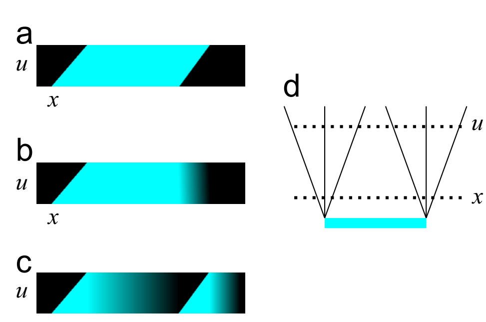

Fig. 7. This toy example provides motivation for the use of gradient magnitude metrics. (a) is an example light field corresponding to a cyan object positioned as in (d) against a black background at some distance behind the x-plane of the two-plane parametrization. (b) shows what a light field generated by our design optimized under a hogel-distance error metric might look like. (c) shows a light field that instead minimizes gradient magnitude difference. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this article.)

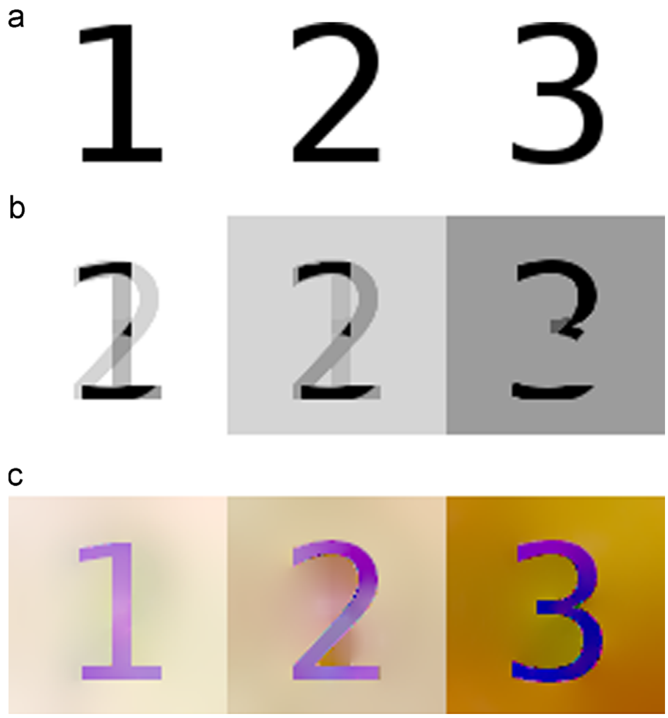

Fig. 8. The basic hogel-distance objective compared with the gradient preserving objective. (a) shows the target light field, with three distinct views, (b) shows a simulation of the result found with a least-squares hogel-distance optimization, and (c) the result using a gradient magnitude preserving objective (in CIE-Lab space).

4.3. Implementation

For both approaches we minimize the appropriate objective function using the SciPy [20] implementation of the L-BFGS-B [21] constrained optimization algorithm.

For the basic hogel-distance case under a 2-norm there is a trivial closed-form expression for the optimization gradient

$$\frac{\partial E}{\partial \mathbf{a}(x,y)} = 2B^T(B\mathbf{a}(x,y) - \mathbf{t}(x,y)) \tag{10}$$

and the effect of each $\mathbf{a}$ is local, so the full optimization of an $80 \times 80 \times 16$ light field takes anywhere from 10 s to a minute depending on content.

For the gradient preserving case at each iteration we estimate the optimization gradient, dE/da, with forward differences. Changing $a_u(x,y)$ affects only nearby terms in x and y, so calculating the optimization gradient can be performed in $O(XYU^2)$ time. In practice U is much smaller than X and Y, so the quadratic term is not a limiting factor.

The optimization gradient calculation is highly parallel, with each $\partial E/\partial a_u(x,y)$ independent, so well suited to a GPU implementation. Our implementation in OpenCL running on a Nvidia GeForce GT 330M GPU takes on the order of 950 ms to calculate the optimization gradient for an $80 \times 80 \times 16$ light field. The number of iterations before convergence, however, is in the thousands, so this approach remains significantly slower than the basic hogel-distance. For all of our results (except the Fig. 8) we used 16 subpixels in between adjacent walls. This is also the number of constrained views.

Initial guess: The objective function (9) has local minima, which our optimization technique can get stuck in. A good initial guess is critical for good results. We have found that a simple linear transformation of the target hogels, forcing them in gamut, works well. In other words we initialize

$$\mathbf{a}(x,y) = \mathbf{t}(x,y) \tag{11}$$

Effectively this takes points originally in the cube gamut of Fig. 4 (a) and shears them into either Fig. 4(b) or (c) depending on the subsequent choice of B.

5. Results and applications

5.1. Experimental setup

For prototype purposes, we constructed a frame consisting of just the vertical occluding walls on an Objet Connex 500 3d printer. Then, using a commodity inkjet photo printer, we print desired textures onto photo paper, and place the occluding frame on top of the paper. This is shown in Fig. 9. The printer material used was very pliable, and had a tendency to warp even under its own weight, hence the horizontal stiffening members. These are not fundamentally required by our design, and future work should investigate the possibility of superior fabrication methods.

We envision for low-cost manufacture printing directly onto heavy card-stock or paper, and using a stamping process to introduce the occluding walls.

We show examples of our system being used for a variety of applications below. In all cases, we chose to include high-resolution images in the paper even though this might lead to aliasing when viewing on a screen or printed copy. Our experience with the physical prototypes is that aliasing is not a problem. Refer to our video results, which show the results much more clearly.

5.2. Application to light fields

We constructed target light fields by rendering a series of shifted views of 3d scenes. For light fields, we use the gradient magnitude preserving optimization described in Section 4.2. Fig. 10(a) shows the result of running the optimization on a light field with fairly simple edge structure, using the L_2 norm for the outer norm of Eq. (9). The white “halo” flips the direction of the gradients for some edges, in order to fit their hogels in gamut.

Fig. 10(b) shows a more complex light field, here of the standard “dragon” model [22]. Here the L_1 norm was used. Note that the halo is much less pronounced in this case.

5.3. Application to animation

Another application of our surfaces is displaying “animation” light fields rather than physical ones. In such a display, each view corresponds to a frame of an animation. This is a standard application of parallax barriers and lenticular displays. In some

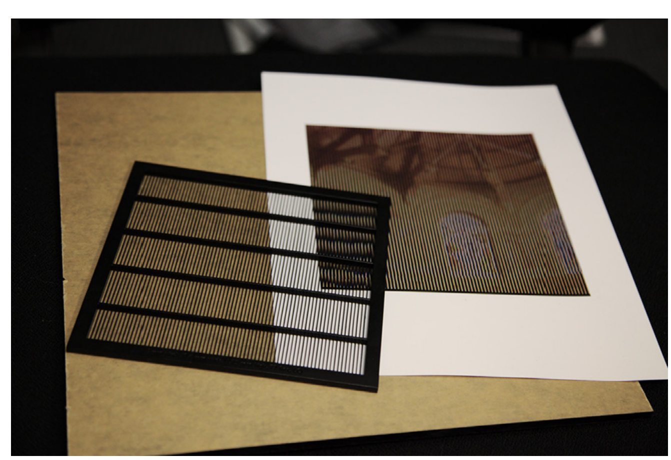

Fig. 9. One of our occlusion frames next to a sample texture. The thick horizontal bars are for stiffening, due to the 3d printed material's tendency to warp.

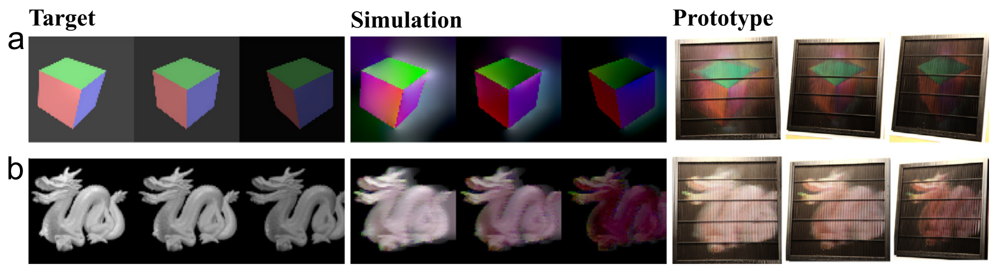

Fig. 10. Light field display with our technique. We show only three views here, however there are 16 total. Leftmost column is target, centre is simulation, and right is physical prototype. Note how in (a) the gradient of the rightmost edge was inverted in order to fit it in gamut. The L_2 norm was used for the outer norm for (a) and L_1 for (b).

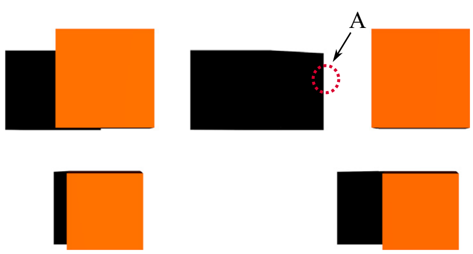

Fig. 11. The problem of overhang in relightable scenes. Here two cubes are relit, the upper cube lies above the ground-plane (so has overhang) and the lower cube does not. The marked edge A creates out-of-gamut hogels. As the light moves further from the centre, points at that edge become brighter rather than darker.

cases this is a harder problem. Basic Lambertian light fields typically either have achievable edges, or ones whose gradient must be flipped (refer to Fig. 7). General animations do not have this property, however for fairly short animations where the views are correlated good results can be achieved.

Fig. 12(a) shows an example of an animation of a growing star that our technique is well-suited for.

5.4. Application to relightable scenes

One nice property of our technique is that there is a duality between lighting direction and viewing direction. Just as a centrally (or ambient) lit surface produces a sequence of views as the viewer position changes, the same sequence of views is produced by keeping the viewer fixed at a central location and shifting a (point) light-source.

This allows our technique to display scenes that change with lighting direction. An application of this is "relightable" scenes where the light source location realistically changes the shading of the scene.

The gradient-magnitude preserving approach of Section 4.2 is poorly suited here: flipping the gradient direction for a shadow edge is perceptually unacceptable, as the edges of natural shadows are of course by definition darker on the shadowed side. However, for Lambertian scenes the hogel-distance approach works well. Recall the hogel gamut is such that the central view (or in this case, lighting direction) is the brightest, with views monotonically darkening as they move off-centre. For Lambertian objects this property almost holds.

Fig. 12(b) shows an example of relighting a face. For a central lighting condition no shadows are visible, and as the light shifts left or right, the shadows grow larger. This, then, is mostly in gamut, as the shadows grow as the light source moves to more extreme angles. This is also shown for a continuously moving light source in our video results.

Faces are in-gamut partly because they lack overhang. Lambertian scenes with overhang are problematic. The "leading" edge of the shadow is in-gamut, however overhang causes a "trailing" edge which is out of gamut. This is shown in Fig. 11. A perceptually motivated shadow-preserving objective function would be needed to allow these scenes to be rendered, if indeed there exist pleasing in-gamut solutions at all.

Our method has some advantages compared to the previous work on passive relightable scenes. Malzbender et al. [9] have shown a setup using a spherical mirror array overlaid with a printed transparency. Parallax walls have a simpler geometry that lets them scale to large setups as in architectural applications. Another advantage of our approach is that for a single light orientation, multiple subpixels are simultaneous lit resulting in a better viewing experience from close distance. In Malzbender et al.'s setup only a single subpixel is lit, being surrounded by black subpixels. Their setup's main advantage is a much more general gamut, but, as our relighting results show, this generality is often not necessary. For instance, see Fig. 13. These scenes all work well, despite being partially out of gamut in every case.

Both our methods also share some limitations. They distribute incoming light in many directions and therefore cannot represent specular surfaces at their full brightness. Ours due to diffuse reflection, theirs due to spatial multiplexing. In addition, both approaches have a preferential position for the viewer in the up direction.

5.5. Application to HDR display

Recently Dong et al. presented a novel technique for displaying high dynamic range (HDR) images, in which altering either the lighting-direction or the viewing direction of a surface displayed different exposures of an HDR scene [19]. Their approach used glossy inks to achieve this view-dependance.

Our design can also be used for these sorts of applications. Note that an HDR exposure stack has the property of monotonically decreasing brightness, consistent with our display basis B_s. So these exposure stacks are, up to noise, in gamut. Simply optimizing under the hogel-distance objective works very well as shown in Fig. 12(c).

5.6. A note on result quality

Although the results presented above clearly show the technique's applicability to a wide range of applications, they are all of fairly low resolution. Part of this is due to being statically reproduced on the page, and we encourage the reader to view our video results. However, further, we were limited in this case by the size and resolution of our 3d printers, as well as the strength of the 3d printed material. These constraints limited overall prototype size, and required fairly thick walls. However with other manufacturing techniques there would be no theoretical obstacle in improving this resolution.

6. Conclusions and future work

We have demonstrated that simple, easily manufactured surfaces can generate surprisingly complex light fields. By characterizing the

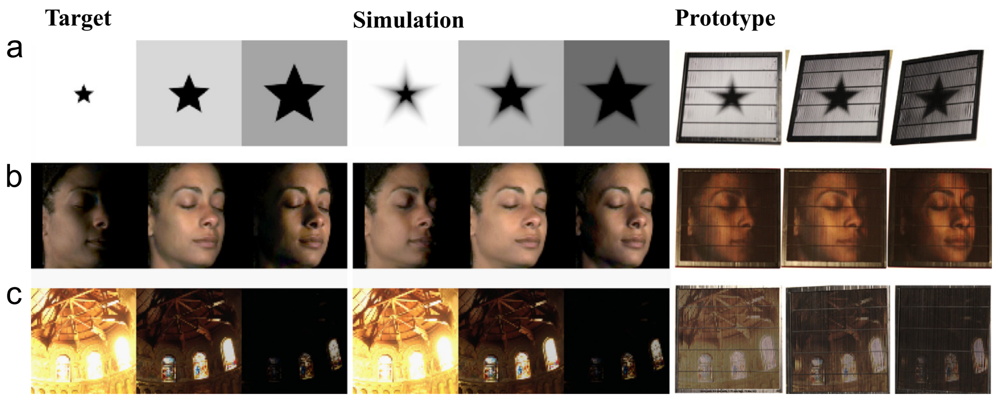

Fig. 12. Other applications. We show three distinct images out of 16 total views in each case. (a) Animation display, (b) Relighting a face. Dataset from Debevec et al. [23] and (c) HDR exposure stack display of Memorial cathedral. Dataset from Debevec et al. [24].

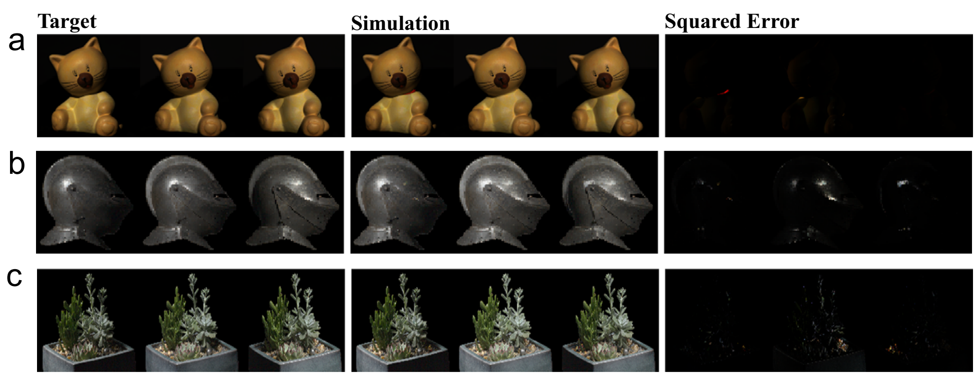

Fig. 13. More relighting results in simulation. Here the third columns show squared error, so bright points indicate error. In general the fits to general reflectance fields are good, with artifacts around non-Lambertian points. Note that these results are brighter than those of Fig. 12(b), because they were only optimized for 3 distinct lighting conditions. As such the maximum brightness of the extremal views in this case are 2/3 the brightness of the central view. (a) Relighting a toy. This object has minimal self-shadowing and is Lambertian, so is mostly in-gamut, (b) Relighting a helmet. Note the error is focused around the specularity, where reflectance field falls out of gamut and (c) Relighting a plant. The error here is focused around the self-shadowing of the plant, but is not perceptually problematic.

light field basis of a surface, and using perceptually motivated gamut-mapping to project into this basis, otherwise apparently difficult results can be achieved.

Limitations: Most arbitrary content is not in the achievable gamut of our surfaces (HDR exposure stacks being a notable exception), as such our gamut mapping necessarily alters desired content. For physical light fields this mapping has the effect of altering the colours while preserving the geometry. Alternatively we could alter the geometry and preserve the colours, but we cannot have both.

Our model treats the surface as uniformly lit, which does not take into account ambient occlusion effects around the sides of the walls. It also assumes perfectly black walls, while in reality diffuse inter-reflections alter wall colour and cause perceptible noise.

Finally, our model does not take into account prefiltering of the input light field to avoid aliasing. Approaches such as those of [25,1] could serve to mitigate these issues.

Future directions: We consider in depth a particular surface geometry in this paper. Other work considers arbitrary surface geometry with fixed albedo [5,6]. Combining these two approaches may lead to better results.

Other non-adaptive geometries may also be worth exploring, for instance Muff [8] examines surfaces with repeating triangular patterns, and Malzbender et al. use repeated concave pits for reflectance field display [9]. Under our analysis these surfaces would simply cause different basis matrices **B**. Perhaps for some classes of light fields, these alternate bases would be preferable. In this work we investigate 3d light fields with only horizontal parallax. Surface geometries including horizontal walls may allow us to display full 4d light fields.

Beyond non-adaptive geometries, adapting the surface-orthogonal geometry to the content of the scene being displayed, as done in the previous work, would be an interesting direction not explored in this work.

A thorough exploration of different perceptually motivated objective functions, similar to our gradient magnitude preserving Eq. (9), might lead to better results, both for the light field case and for the relighting case where that equation is not appropriate, as discussed in Section 5.4.

Finally, although in this work our prototypes were built with a 3d printer (see Section 5), we envision two future directions: building them on single sheets of paper, stamped with our height field structure, or building them at extremely large scales.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version, at http://dx.doi.org/10.1016/j.cag.2013.07.002.

References

[1] 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 on Graphics 2011;30(4).
[2] Nayar SK, Belhumeur PN, Boult TE. Lighting sensitive display. ACM Transactions on Graphics 2004;23(4):963–79.
[3] Weyrich T, Peers P, Matusik W, Rusinkiewicz S. Fabricating microgeometry for custom surface reflectance. ACM Transactions on Graphics 2009;28(3):32.

[4] Alexa M, Matusik W. Images from self-occlusion. In: Proceedings of the international symposium on computational aesthetics in graphics, visualization, and imaging. CAe'11. New York, NY, USA: ACM; 2011. p. 17–24.

[5] Alexa M, Matusik W. Reliefs as images. ACM Transactions on Graphics 2010;29(4):60.

[6] Bermano A, Baran I, Alexa M, Matusk W. Shadowpix: multiple images from self shadowing. Computer Graphics Forum 2012;31(2pt3):593–602.

[7] Wu H, Dorsey J, Rushmeier H. Physically-based interactive bi-scale material design. In: Proceedings of the 2011 SIGGRAPH Asia conference. SA '11. New York, NY, USA: ACM; 2011. p. 145:1–10.

[8] Muff S. Stereoscopic surfaces. Master's thesis. ETH Zurich; 2011.

[9] Malzbender T, Samadani R, Scher S, Crume A, Dunn D, Davis J. Printing reflectance functions. ACM Transactions on Graphics 2012;31(3):20:1–20:11.

[10] Lanman D, Hirsch M, Kim Y, Raskar R. Content-adaptive parallax barriers: optimizing dual-layer 3D displays using low-rank light field factorization. ACM Transactions on Graphics 2010;29(6):23.

[11] Baran I, Keller P, Bradley D, Coros S, Jarosz W, Nowrouzezahrai D, et al. Manufacturing layered attenuators for multiple prescribed shadow images. Computer Graphics Forum 2012;31(2pt3):603–10.

[12] Wetzstein G, Lanman D, Gutierrez D, Hirsch M. Computational displays: combining optical fabrication, computational processing, and perceptual tricks to build the displays of the future. In: ACM SIGGRAPH 2012 Courses. In: SIGGRAPH '12. New York, NY, USA: ACM; 2012. p. 4:1–159, ISBN 978-1-4503-1678-1. http://dx.doi.org/10.1145/2343483.2343487.

[13] Mitra NJ, Pauly M. Shadow art. ACM Transactions on Graphics 2009;28(5):156.

[14] Fattal R, Lischinski D, Werman M. Gradient domain high dynamic range compression. In: Proceedings of the 29th annual conference on computer graphics and interactive techniques. SIGGRAPH '02. New York, NY, USA: ACM; 2002. p. 249–56.

[15] Weyrich T, Deng J, Barnes C, Rusinkiewicz S, Finkelstein A. Digital bas-relief from 3d scenes. In: ACM SIGGRAPH 2007 papers. SIGGRAPH '07. New York, NY, USA: ACM; 2007.

[16] Levoy M, Hanrahan P. Light field rendering. In: SIGGRAPH '96. New York, NY, USA: ACM; 1996. p. 31–42.

[17] Gortler SJ, Grzeszczuk R, Szeliski R, Cohen MF. The lumigraph. In: Proceedings of the 23rd annual conference on computer graphics and interactive techniques. SIGGRAPH '96. New York, NY, USA: ACM; 1996. p. 43–54.

[18] Lucente M. Diffraction-specific fringe computation for electro-holography. PhD thesis. MIT; 1994.

[19] Dong Y, Tong X, Pellacini F, Guo B. Printing spatially-varying reflectance for reproducing hdr images. ACM Transactions on Graphics 2012;31(4):40.

[20] Jones E, Oliphant T, Peterson P, et al. SciPy: Open source scientific tools for Python. 2001. URL ⟨http://www.scipy.org/⟩.

[21] Byrd RH, Lu P, Nocedal J. A limited memory algorithm for bound constrained optimization. SIAM Journal on Scientific and Statistical Computing 1995;16(5):1190–208.

[22] Curless B, Levoy M. A volumetric method for building complex models from range images. In: SIGGRAPH '96. New York, NY, USA: ACM; 1996. p. 303–12.

[23] Debevec P, Hawkins T, Tchou C, Duiker HP, Sarokin W, Sagar M. Acquiring the reflectance field of a human face. In: SIGGRAPH '00. New York, NY, USA: ACM Press, Addison-Wesley Publishing Co.; 2000. p. 145–56.

[24] Debevec PE, Malik J. Recovering high dynamic range radiance maps from photographs. In: ACM SIGGRAPH 2008 classes. SIGGRAPH '08. New York, NY, USA: ACM; 2008. p. 31:1–10.

[25] Bruneton E, Neyret F. A survey of nonlinear prefiltering methods for efficient and accurate surface shading. IEEE Transactions on Visualization and Computer Graphics 2012;18(2):242–60, http://dx.doi.org/10.1109/TVCG.2011.81.