

Superresolution Reflectance Fields: Synthesizing images for intermediate light directions

Martin Fuchs¹, Hendrik P. A. Lensch¹, Volker Blanz², and Hans-Peter Seidel¹

¹MPI Informatik, Saarbrücken, Germany ² Universität Siegen, Germany

Abstract

Captured reflectance fields tend to provide a relatively coarse sampling of the incident light directions. As a result, sharp illumination features, such as highlights or shadow boundaries, are poorly reconstructed during relighting; highlights are disconnected, and shadows show banding artefacts. In this paper, we propose a novel interpolation technique for 4D reflectance fields that reconstructs plausible images even for non-observed light directions. Given a sparsely sampled reflectance field, we can effectively synthesize images as they would have been obtained from denser sampling. The processing pipeline consists of three steps: (1) segmentation of regions where intermediate lighting cannot be obtained by blending, (2) appropriate flow algorithms for highlights and shadows, plus (3) a final reconstruction technique that uses image-based priors to faithfully correct errors that might be introduced by the segmentation or flow step. The algorithm reliably reproduces scenes that contain specular highlights, interreflections, shadows or caustics.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer graphics]: Three-Dimensional Graphics and Realism I.3.7 [Computer graphics]: Color, shading, shadowing and texture

Keywords: Reflectance Fields, Image-Based Lighting, Upsampling

1. Introduction

Reflectance fields have been shown to be a powerful approach for creating photo-realistic images of objects or scenes in new illumination conditions. The key idea of reflectance fields, and the basis for almost all work in this field, is the observation that, due to the principle of superposition of light, the combination of multiple light sources can be simulated by adding the pixel values of images taken with separate light sources. The effect of an environment map on a scene is simulated by a weighted sum of images taken at different illuminations, such as point lights at different directions of a light stage setup.

For high-quality results, however, reflectance field methods have been quite costly in terms of acquisition time and storage space since they require one input image for each incident light direction. Faithful reproduction of high frequency illumination effects such as sharp shadow boundaries, specular highlights or caustics require a very dense sampling of the light directions. If the spacing between illumination directions is too large, salient artifacts occur in

simulations of extended light sources (see Figure 1): In specular highlights on glossy surfaces, the original sampling pattern is clearly visible as a pattern of spatially separated highlights. In simulations of moving directional lights, the synthesized sequence will show aliasing artifacts: Adjacent highlights appear and disappear, rather than moving smoothly over the reflective surface. Similar aliasing effects can be observed in moving shadows. In general, smooth results can only be expected if the sampling density is appropriate for the angular resolution of the illuminations in terms of the sampling theorem [FLBSar]. For light stages with LEDs or small flashlights, dense sampling may well require tens of thousands of images, which are too costly to acquire.

We therefore propose a novel interpolation and upsampling scheme that takes as input a sparsely sampled reflectance field captured with a point light source and produces a plausible reflectance field with much higher resolution, for instance, supersampling it from 230 captured to about 4000 synthesized illumination directions. With the constructed superresolution reflectance field, the motion and




Figure 1: Relighting a sparsely sampled reflectance field (230 input images) with a small area light source. Top: Linear interpolation of the input samples results in banded shadows and disconnected highlights. Bottom: Our novel interpolation technique upsamples the same input reflectance field, effectively increasing its resolution to 3547 samples, and is capable of plausibly reproducing smooth shadows, shadows of semi-transparent objects, highlights and caustics.

appearance of high frequency illumination effects can be approximated without requiring to capture thousands of samples.

The intermediate reflectance images are created by first separating the input images into regions where simple linear interpolation is sufficient, and regions where it is not. The algorithm distinguishes between highlights, shadow contours and regions that change slowly with illumination, such as diffusely lit surfaces and the inner regions of shadows. For the components that need to be explicitly moved, we apply appropriate flow algorithms. In order to predict the appearance of the scene for an intermediate light direction we finally compose the warped components and regularize the solution using image-based priors. As we separate the prediction whether a pixel should be shadowed or not from the estimation of its appearance when actually shadowed, we maintain local texture detail while reconstructing a sharp shadow boundary.

The benefit of our approach is that from low sampling rates with high-spatial frequency illumination, we can simulate intermediate illuminations both at high spatial frequencies, such as point lights, and simulate extended light sources, such as those found in many light fields. From a theoretical perspective, the algorithm taps new sources of

information about reflectance from sparse samples by modeling and separating the tight connection between the effects in the image plane and in the angular domain.

2. Related Work

Most reflectance field capturing approaches sample the light source direction sparsely by moving a point light source [DHT*00, MDA02, MPZ*02] or a projector [MPDW03, SCG*05, GTLL06] to a discrete set of positions on the sphere around the object. High frequency effects in the light domain, e.g. shadow, highlights or caustics, require a much higher sampling density. This can be obtained by literally capturing thousands of samples with a dual light stage [HED05] at immense acquisition costs, or by trying to parallelize the acquisition of multiple directional samples using an adaptive scheme [SCG*05] or a stochastic approach based on wavelet noise [PD05]. While considerably faster, the latter ones have the disadvantage that dim features in the reflectance field might be captured at high noise levels because they are always observed in conjunction with some potentially much brighter reflections.

In this paper, we instead propose an advanced non-linear interpolation technique that can correctly reproduce the object's appearance for in-between light source positions from a sparse sampling. Simple blending between input samples is typically not sufficient.

One way to estimate intermediate samples is to approximate the reflectance field locally by an explicit model. Various models have been proposed to describe the apparent BRDF for the surface point visible at a given camera pixel, such as analytic BRDF models [DHT*00], which, however, do not account for global effects. More general functions such as spherical harmonics and wavelets [MPDW04] on the other hand can at best provide a smooth interpolation between given samples. High frequencies will not be introduced to pixels which did not observe them. In order to increase the sampling density locally the reflectance sharing approach combines samples from multiple surface points [LLSS03, ZREB06]. This technique has so far been demonstrated for opaque surfaces of known geometry only. In this paper we present interpolation techniques for reflectance fields of non-opaque surfaces and unknown geometry.

The problem of interpolation between light source directions is related to the problem of view interpolation, for which various techniques have been proposed, e.g. optical flow [BA96, BBPWk04] or level set blending [Whi00]. Special solutions to handle illumination effects include flow for specular surfaces [RB06] or interpolation for specularly refracting materials [MPZ*02]. However, applying view interpolation techniques to the input images of a reflectance field directly has the inherent problem that the scene itself and all its texture is static while the illumination creates apparent motion on top of the static structures. Barrow et al. [BT78] therefore proposed the separation into intrinsic

and illumination images. For reflectance fields of mostly diffuse scenes, Matsushita et al. [MKL*04] used the illumination image to detect shadows which are then moved according to the position of the light source with the help of an explicit 3D model of the scene. Applying view-interpolation techniques on thresholded binary images of detected shadow regions, Chen and Lensch [CL05] generated smoothly moving shadow regions. They worked on 6D reflectance fields where the illumination is controlled by a projector. This allowed them to turn off the direct illumination to the shadow regions while still considering the indirect illumination from other scene parts.

Our upsampled reflectance fields are regularized by applying image-based priors introduced by Fitzgibbon et al. [AF03]. They performed view interpolation and implicitly reconstructed a depth map subject to the constraint that the interpolated view is locally consistent with the recorded image data. The constraint enforces that any pixel neighborhood in the interpolated view has to occur somewhere in the input views. Wexler et al. [WSI04] further extended this idea to fill in holes in a space-time video cube. Without any additional information, relying on image-based priors only, the holes are filled in by a multi-resolution framework using 3D neighborhoods. At each level the solution of the previous level is regularized to match the image-based priors at that resolution. Upsampling reflectance fields can be viewed as filling holes in a 4D structure. In contrast to the work by Wexler et al., our input data unfortunately provides 2D slices only, i.e. priors for 2D neighborhoods. We therefore need to include the reconstruction step based on highlight and shadow maps in order to recover 4D consistency.

3. Overview

Reflectance Fields represent a powerful tool to describe light transport in real world scenes. In the 4D case, the reflectance field is defined as a function $R(x, y, \theta, \phi)$ that maps the distant light coming from the direction (θ, ϕ) in polar coordinates to the brightnesses $R(x, y, \theta, \phi)$ at pixel positions (x, y) in an image of $w \times h$ pixels. Color channels are treated separately.

For a given environment map $L(\theta, \phi)$, images can be rendered according to the Reflectance Integral for Image-Based Relighting:

$$I(x, y) = \int_{\theta \in [0; \pi]} \int_{\phi \in [0; 2\pi]} L(\theta, \phi) \cdot R(x, y, \theta, \phi) d\phi \sin \theta d\theta \tag{1}$$

Captured reflectance fields usually consist of a number of n slices $R_i(x, y)$, corresponding to images capturing the appearance of the scene for a single illumination direction (θ_i, ϕ_i) . In order to render these with continuous environment maps, interpolation along the θ, ϕ domain is necessary [MPDW04]. Unfortunately, such a direct interpolation may be insufficient if n is small; however smooth the interpolation may be, it does not ensure that high-frequency features like shadow boundaries or highlights vary smoothly in the image domain.




Figure 2: The incident light directions of a sparsely sampled ($n = 230$) reflectance field visualized as black dots, and a connecting Delaunay triangulation which defines barycentric coordinates on which linear interpolation can be performed. We increase the resolution by synthesizing samples for in-between light directions (red dots).

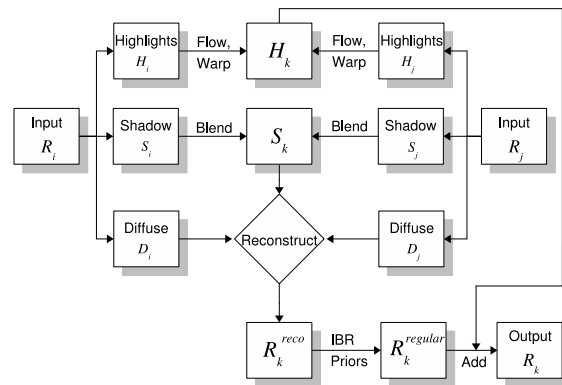


Figure 3: Input slices R_i, R_j of the reflectance field for incident light angles θ_i, ϕ_i and θ_j, ϕ_j , respectively, are separated into highlights, shadows and a diffuse component. We interpolate the components separately and combine them to a new slice R_k for an intermediate light direction.

In our approach, we reduce the interpolation problem by synthesizing plausible intermediate images R_k at a much higher density in the domain of light directions (see Figure 2) and inserting them into the measured field. From there, we can proceed with any interpolation scheme; for this paper, we show results based on linear interpolation in barycentric coordinates on a Delaunay triangulation of the input images. We generate the intermediate images R_k by processing the input images R_i according to the pipeline depicted in Figure 3. The pipeline consists of three main stages: segmentation and labeling of the input data, separate upsampling of the highlight and shadowmaps, and reconstruction at the target resolution subject to image-based priors.

The segmentation stage (Section 4) separates out highlights and shadows which are treated separately in our pipeline. From the input images R_i we extract the images H_i , that just contain the highlights, and reconstruct images D_i where the information beneath the highlights is plausibly reconstructed using image-based priors. In a second

step, a binary shadow map S_i is computed, which encodes whether a pixel is likely to be in shadow ($S_i(x, y) = 0$), or not ($S_i(x, y) = 1$). The shadow labeling does not need to be very precise since it is used mainly to obtain the correct movement of shadow boundaries, but not the interior intensity.

In the second stage of the pipeline (Section 4), we upsample the highlight and shadow data to higher resolutions in the θ, ϕ domain: we apply optical flow to warp the highlight images H_i and perform level-set image blending [Whi00] on the shadow maps S_i .

In the third stage, a tentative intermediate image R_k^{reco} is synthesized. From the warped shadow images we determine for each pixel how to interpolate its reflectance from shadowed or lit samples. We then apply image-based priors in order to regularize our solution, removing intensity or high frequency artefacts introduced by incorrect segmentation. Finally, we add the warped highlight layer and arrive at the interpolated image R_k .

4. Segmentation

Our method begins with the extraction of semantic maps of highlights and shadows for our input images R_i . In order to obtain a criterion to determine the status of each pixel, we sort the measured reflectance values $R_i(x, y)$ for each pixel, resulting in images P_i , where P_0 contains the dimmest reflectance for each pixel and P_{n-1} the brightest. From these pictures we select a picture P where no pixel is a highlight nor a shadow. For an input reflectance field consisting of $n = 256$ images, we found $P = P_{230}$ to be a robust choice. Based on P , shadow and highlight maps are generated in the following steps.

4.1. Specular Highlights

The separation of highlights is done separately for each input image R_i in several steps. At first, a binary map $M_i^{high}(x, y)$ encodes whether a pixel $R_i(x, y)$ contains highlight information or not. Initially, $M_i^{high}(x, y)$ is 1 iff $R_i(x, y) \geq f_{highlight} \cdot P(x, y)$, where $f_{highlight}$ is a constant factor for all images; in our experience, $f_{highlight} = 2$ or 4 is a good choice. Then, we dilate $M_i^{high}(x, y)$ with a 3×3 or 5×5 box filter, growing the highlight regions to include possible haze around them.

Based on M_i^{high} , we generate the highlight-free images D_i by removing all possible highlight pixels from R_i and applying the hole filling algorithm by Wexler et al. [WSI04] (see Section 6.2) to reconstruct plausible texture information replacing the highlights. As texture references, we use R_i and the images for neighboring light configurations, but we remove potential highlight pixels. In modification to Wexler et al.'s original algorithm, we do not perform an outlier detection, and add a term to the region lookup which favors close-by regions, as this improves the performance of the lookup structure.

Finally, the highlight map $H_i(x, y) = R_i(x, y) - D_i(x, y)$ is computed as the per-pixel difference. Figure 4 illustrates the highlight detection and removal process.

4.2. Shadow Boundaries

In contrast to the highlight maps, the shadow maps are generated not independently for each input image, but independently for each pixel at position (x, y) . We perform a region growing on lit regions and treat the remainder as shadow.

More precisely, a shadow map $S_i(x, y)$ is generated for all directions (θ_i, ϕ_i) , with initially $S_i(x, y) = 1$ for “lit” iff $R_i(x, y) \geq P(x, y)$, and $S_i(x, y) = 0$, otherwise. In order to mitigate camera noise one can perform this segmentation based on a prefiltered version of R_i . We then perform the following iterative update: Let (θ_i, ϕ_i) be the direction for which pixel (x, y) is currently labeled as shadowed ($S_i(x, y) = 0$), and (θ_j, ϕ_j) a neighboring direction for which the pixel has been labeled as lit ($S_j(x, y) = 1$). For image R_i , the pixel is re-labeled as lit ($S_i(x, y) := 1$) if the following criterion is met: $R_i(x, y) > f_{lit} \cdot R_j(x, y)$. We usually chose $0.66 \leq f_{lit} \leq 0.9$. The region growing is iterated until no more updates take place. Thus, the lit regions grow until the difference to neighboring directions becomes so large that linear blending is likely to fail, indicating a shadow boundary at this location.

Figure 5 illustrates the detection of shadow maps. Note that while the results may be noisy in the (x, y) domain, the scheme is continuous in the (θ, ϕ) domain, where the reconstruction later takes place.

5. Upsampling of highlight and shadow data

Before we can synthesize reflectance images I_k at the full resolution, we need to procure estimations of the shadow and highlight distributions for this resolution. Let S_i and S_j be shadow maps, and H_i and H_j be highlight maps for two input images R_i, R_j that have adjacent light directions (θ_i, ϕ_i) and (θ_j, ϕ_j) in the mesh provided in Figure 2. In this section, we will explain how to generate intermediate maps S_k and H_k for the in-between direction. The full target resolution is then later obtained by iteratively subdividing the mesh edges and recomputing the triangulation until the desired resolution is achieved, i.e. information is available at all red dots in Figure 2.

5.1. Upsampling Specular Highlights

The generation of H_k is rather straight-forward. We apply the optical flow algorithm by Brox et al. [BBPWk04] both to obtain a flow field from H_i towards H_j , and from H_j to H_i . We clamp the HDR data of the highlights to obtain a smaller dynamic range, which makes matching easier. Using the flow fields, we warp H_i and H_j with full dynamic range to the halfway position, and blend the result linearly to obtain H_k . The flow algorithm is capable of generating very smooth fields, which is useful for our application, as it allows to drag highlight data along even if only neighboring parts matched. As we segment our highlights carefully and the separated highlight image is mostly black, this does not introduce artefacts.

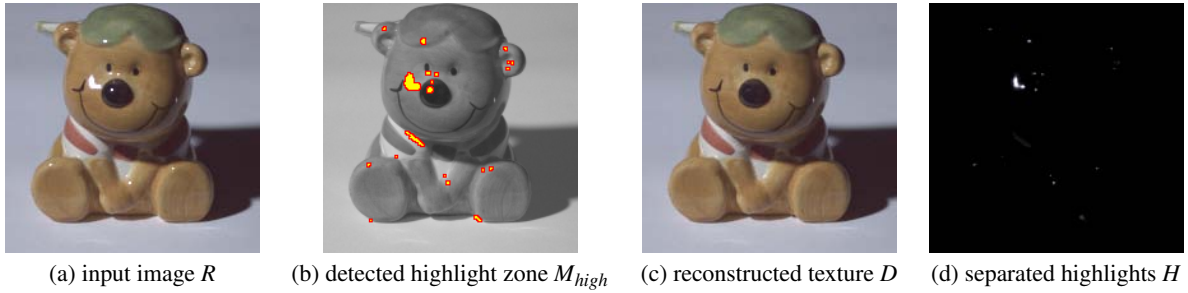


Figure 4: Separation of highlights. By intensity analysis in the input image R (a), a region of definite highlights M_{high} is detected (b, colored in yellow) which is dilated (b, colored in red) for robustness. Texture reconstruction allows to estimate the appearance D below the highlights (c). By subtracting D from R , the highlight layer H for this image is estimated (d). Note that the reconstructed highlights have smooth boundaries.




Figure 5: Computation of shadow maps. To the left, the top row shows a plot of acquired reflectance in polar parameterization for three pixels (to the right of the figure's left foot). Below, the segmentation in shadow (black) and lit (white) regions is displayed. The image to the right shows a slice of the shadow map in image space for a fixed incident light angle (θ, ϕ) . While noisy in image space, the shadow computation is stable along the (θ, ϕ) directions.

5.2. Upsampling shadow data

For the interpolation of the shadow maps S_i , we employ the level-set blending approach of Whitaker [Whi00], which generates intermediate shapes of level sets within images. In contrast to optical flow, this algorithm is more tolerant to noise in the image domain, estimating the correct movement of the shadow boundaries while keeping fixed structures constant. In contrast to the original paper, we replace the up-wind scheme for the computation of partial derivatives with simple central differences, as this yields more stable results for our application. Since the level-set blending algorithm produces continuous images even for sharply segmented input data 0 or 1, we re-quantize the output again. This produces plausible shadow boundaries for in-between images (see the central row of Figure 6).

6. Image Synthesis

Having generated the S_k and H_k for the target resolution, we can now synthesize the output images R_k . Let (θ_k, ϕ_k) be the light direction for which the synthesis is required, and (θ_a, ϕ_a) , (θ_b, ϕ_b) , and (θ_c, ϕ_c) be the triangle of observed input directions from Figure 2 which contains (θ_k, ϕ_k) . The reconstruction takes three steps: first, reconstruct the diffuse appearance R_k^{reco} using the estimated shadow maps, second, regularize the estimation using IBR priors, yielding $R_k^{regular}$, and third, add the highlight layer H_k resulting in the reconstruction R_k . Figure 6 illustrates the different reconstruction

steps. In the figure, the highlight map H_k has already been added to the reconstruction results based on the shadow map.

6.1. Reconstruction Based on Shadow Maps

We first reconstruct the initial diffuse appearance R_k^{reco} , given the upsampled shadow maps S . For the construction of R_k^{reco} , we need to obtain estimates for every pixel's appearance both for the hypothesis R_k^{lit} that the pixel may be lit, and for the hypothesis $R_k^{shadowed}$ that it is in shadow. In this step, the shadow maps serve two purposes: a prediction whether a given pixel (x, y) should be shadowed or lit, and a statement whether the input images $R_{a,b,c}(x, y)$ describe the pixel in lit or shadowed state. However, as the shadow maps are binary and possibly imprecise, they cannot be used directly. Instead, we compute a blurred shadow map \bar{S}_k by averaging over the 10 nearest light directions, which we use as smooth approximation of the relative shadowedness in pixel $R_k(x, y)$. Analogously, we also compute smoothed shadow maps $\bar{S}_{a,b,c}$, which are used for estimating the shadow state of $R_{a,b,c}(x, y)$; optionally, one can average over more neighbors in order to be more conservative (see Figure 7). We can now compute the values $R_k^{shadowed}(x, y)$ and R_k^{lit} for the appearance of $R_k(x, y)$ both in shadowed state and in lit state. A set of cases is possible:

1. $\bar{S}_a(x, y) = \bar{S}_b(x, y) = \bar{S}_c(x, y) = 1$, i.e. all neighboring input images have reliably observed the pixel while it

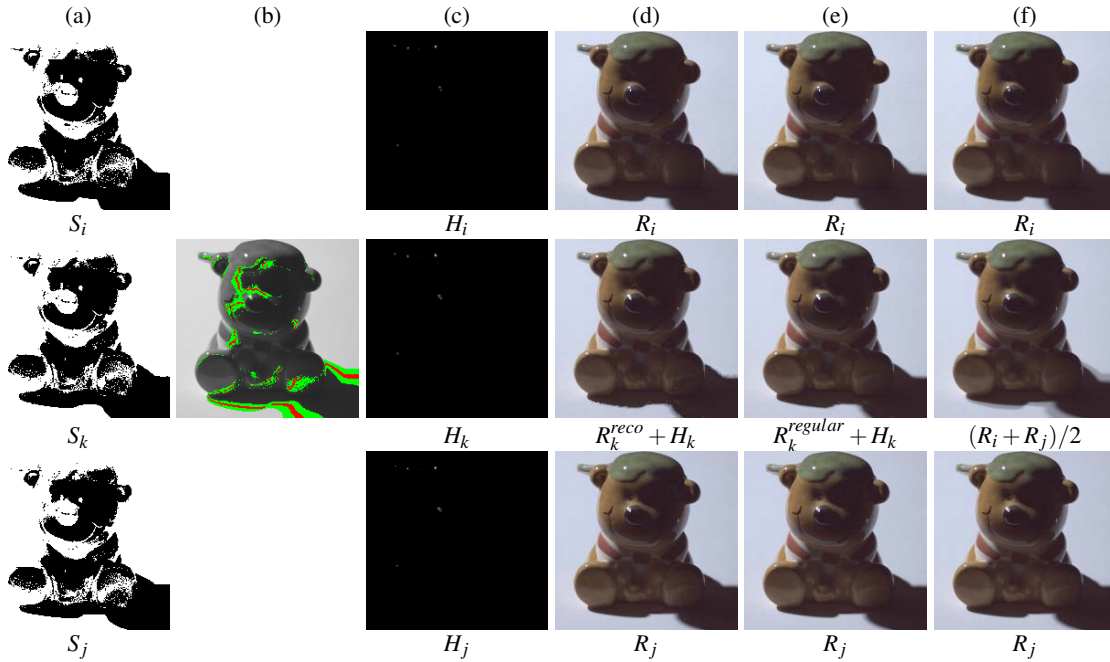


Figure 6: Illustration of the reconstruction process. The bottom and top row contain information from the initial, sparse input field. The in-between row is synthesized. By analyzing the shadow maps (a) we can obtain areas (b) where we either can locally interpolate (gray) or need to locally extrapolate (green and red). The green area is predicted to be fully lit or fully shadowed, the red area is the estimated location of the shadow boundary. After adding the highlight map (c), reconstruction yields a rough picture (d) which is improved using image-based priors (e). Especially the shadow boundary is much more pleasing than in a direct linear interpolation (f).




Figure 7: Polar plot of all blurred shadow maps at a fixed pixel position (white is lit, black shadowed, gray uncertain). The reflectance at light direction k is reconstructed according to the locally estimated shadow state \bar{S}_k by extrapolation from nearby fully lit (white) and fully shadowed (black) measured samples. If the shadow state of all observed conditions a, b, c agrees, direct linear interpolation takes place.

- was lit: We compute $R_k^{lit}(x, y)$ using a direct interpolation using barycentric coordinates α, β, γ as $R_k^{lit}(x, y) = \alpha R_a(x, y) + \beta R_b(x, y) + \gamma R_c(x, y)$.
- $\bar{S}_a(x, y) = \bar{S}_b(x, y) = \bar{S}_c(x, y) = 0$: Similarly, we compute $R_k^{shadowed}(x, y) = \alpha R_a(x, y) + \beta R_b(x, y) + \gamma R_c(x, y)$.
 - For all other cases, we extrapolate a value for $R_k^{lit}(x, y)$ (and $R_k^{shadowed}(x, y)$) by robustly fitting a linear spherical harmonic reconstruction using two SH bands to the observations of nearby light directions, for which a lit state (shadowed state) have been observed.

The output pixel is finally estimated as $R_k^{reco}(x, y) = \bar{S}_k(x, y) \cdot R_k^{lit}(x, y) + (1 - \bar{S}_k(x, y)) \cdot R_k^{shadowed}(x, y)$

Note that in cases where we reconstruct either a fully lit pixel from fully lit neighbors, or a fully shadowed pixel from fully shadowed neighbors, we effectively perform a direct linear interpolation of the pixel value. This implies that outlier pixels, which are wrongly estimated to be shadowed or lit, but are consistently misestimated for a set of close neighbors, tend to be robustly linearly interpolated as well.

6.2. Regularization Using Image-Based Priors

The reconstruction explained above provides an estimate for the output image, but this estimate may be unreliable, e.g. due to misregistrations of shadow or interpolation defects. The affected areas consist of pixels for which a pixel value was computed from extrapolated reflectance data, or in the vicinity of shadow boundaries in $\bar{S}_k(x, y)$. For these, we run a regularization algorithm, which we will describe in more detail below.

We have adapted the approach by Wexler et al. [WSI04] to perform hole filling and regularization on 2D slices of a reflectance field: The algorithm iteratively updates all pixels within the hole region to better fulfill the given priors:

Let E_0 be the initial estimate image a series of iterative updates E_i is constructed. For each window $\mathcal{W}^{E_i}(p, q)$ centered around pixel (p, q) we search for the best matching window $\mathcal{W}^{I_k}(s, t)$ centered around some other pixel in the entire set of input images I_k based on the modified distance measure

$$d(\mathcal{W}^{E_i}(p, q), \mathcal{W}^{I_k}(s, t)) = \arg \min_{\lambda} \sum_{u=-L}^L \sum_{v=-L}^L \|E_i(p+u, q+v) - \lambda I_k(s+u, t+v)\|^2, \quad (2)$$

where L is the radius of the window. In contrast to the original approach we search for the best matching window up to some intensity scale λ . Since our input data is relatively sparse the extension is necessary to find appropriate matches for regions which change their intensity slowly due to the cosine fall-off under diffuse illumination. An efficient lookup can still be achieved by normalizing the input window to unit intensity and then searching for the nearest neighbor using a KD-tree which is built out of normalized windows of the input images. The distance measure is furthermore translated into a similarity measure $s(W^{E_i}(p, q), W^{I_k}(s, t)) = e^{-\frac{d(W^{E_i}(p, q), W^{I_k}(s, t))}{2\sigma^2}}$ which is used in the following to update all pixels (x, y) within the window $W^{E_i}(p, q)$.

Let W_j^E be the set of all windows that contain (p, q) . Each corresponding matched window λW_j^I yields an estimate c_j on what the final color for (p, q) should be. Using the corresponding similarity measures s_j the update is computed as the weighted average:

$$E_{i+1}(p, q) = \frac{\sum_j s_j c_j}{\sum_j s_j}, \text{ compare [WSI04]} \quad (3)$$

Usually, three iterations of this algorithm are sufficient, and using only input images for neighboring light conditions is sufficient. The quality improvement is visible in Figures 8 and 9. Note that in most cases the image-based priors remove high frequency noise but do not corrupt apparent motions which is important to achieve temporal consistency.

In some cases however, the shape of some shadow edges cannot be matched to any input image due to a missing sample with similar orientation. This can lead to slightly deformed shadow edges, as in the slightly too flat shadow of the shiny sphere in the rightmost image in the second row of Figure 9. One approach to circumvent this problem would be to include rotated neighborhoods in the nearest neighbor search.

Once the reconstruction is complete, the warped highlight layer H_k is simply added to the result, thus obtaining the final reconstruction.

7. Results

In order to demonstrate the performance of the proposed method we constructed superresolution reflectance fields for two different scenes: a ceramic bear with highly specular coating (256 input images) and a second, more complicated scene composed of a set of spheres with drastically different reflection and transmission properties, as well as a



Figure 8: Reconstruction results using linear interpolation from 256 images (top row), our upsampling scheme applied three times (center row) compared to a reference rendering with 10 000 input images (bottom row). Left column: simulated extended light source with 15 degrees radius, right column: renderings with a directional source. The linear interpolation produces shadow banding and significant gaps in the highlights. Our upsampling scheme improves significantly on the visual quality.

champagne glass (230 input images). In both cases, the illumination directions were spread equidistantly over the upper hemisphere. Iterating the subdivision process two times, we synthesize about 4000 intermediate positions. The bear data set has a resolution of 200×196 and the sphere set of 453×211 pixels. Timings for a contemporary PC are available in Table 1. The most expensive part of the algorithm is the application of the texture inpainting for the highlight segmentation and the regularization of the output. The table lists the runtimes split into separable units. While I/O operations are considerable if the steps are split this way, it is easy to exploit parallelization and run independent operations for several images concurrently, e.g. on a cluster of PCs. Due to the smaller resolution, the runtime required for the bear data set is considerable lower. For the renderings in Figure 8, we have only generated samples needed for the light situation, and could therefore afford to create a supersampling of three subdivision steps.

In Figure 8 we synthesized the superresolution reflectance field and compare it to the result one would obtain for lin-

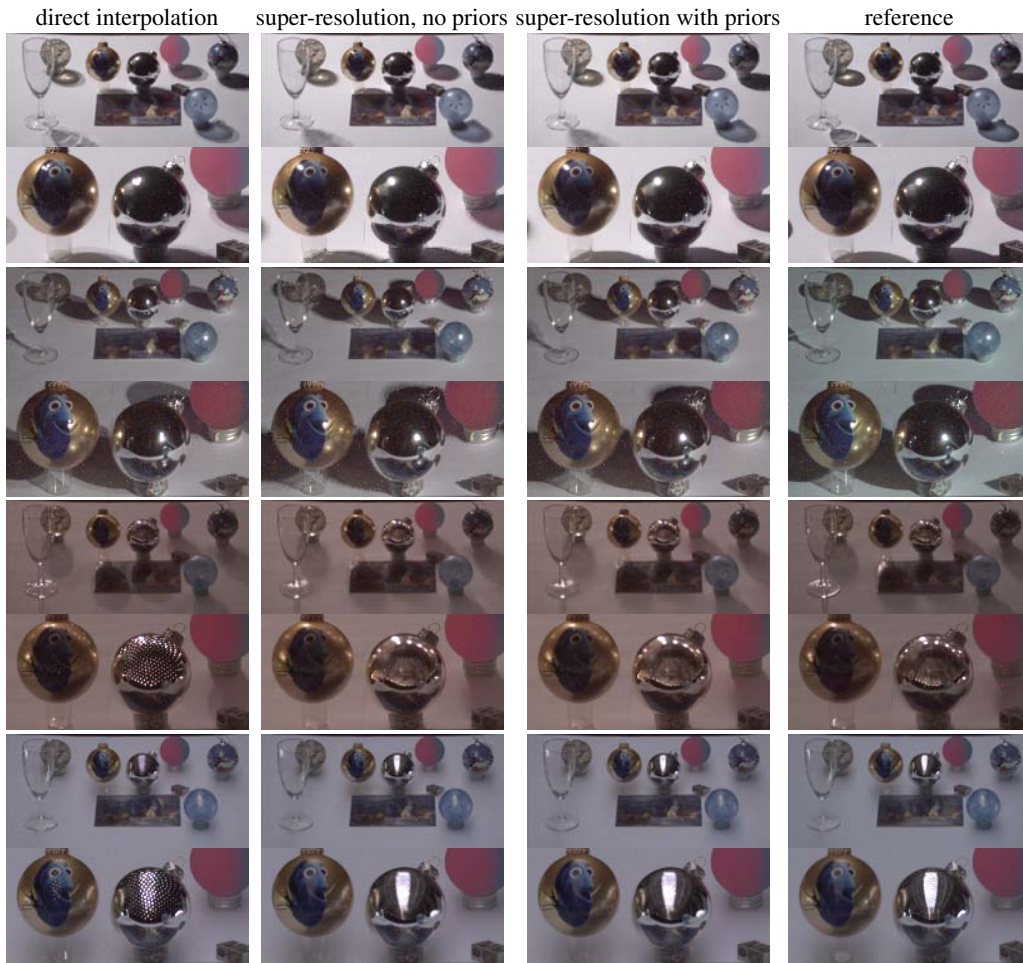


Figure 9: Results of a complex scene with detail zoom-ins for renderings in two point light conditions (upper four rows) and in the Galileo's Tomb and Uffizi environment maps [Deb98]. The super-resolution reflectance field is upsampled from 230 to 3547 images, the reference has 14116 images. Our upsampling approach produces superior results compare to linear interpolation. Applying the regularization (second column from the right) further improves shadow boundaries; for real-world illuminations, the non-regularized version is already close in quality. Note that the reference exposes aliasing which the super-resolution images do not show due to sub-pixel blurring during the upsampling steps.

action	avg. time	
initial sort for threshold image	53 sec.	t
generation of the initial shadow map S_i	62 sec.	t
separation into highlight H_i and D_i	129 sec.	i
level-set blending between two shadow maps	23 sec.	o
optical flow and warp for highlight map	9 sec.	o
blurring of all full resolution shadow maps	7 min.	t
reconstruction of one image without priors	50 sec.	o
reconstruction of one image with priors	< 15 min.	o

Table 1: Timings for the data set in Figure 9. Processing time either for (t)otal field, per (i)nput or (o)utput image.

ear blending and to a reflectance field captured at the full resolution. In the right column, we illuminated the scene from a single intermediate light direction. In the locally, linearly interpolated image double/triple contours from the

three neighboring input images are visible which are removed by our upsampling scheme. While the reconstructed shadow boundary does not perfectly match the reference image, it is still plausible and provides a sharp contour. The left column shows the result after integrating over an extended light source. The linear blending results in a set of clearly separated highlights while our reconstruction shows smooth and connected reflections of the light source. There, the difference to the ground truth is only marginal. Not only the highlights, but also the cast shadows in the scene are reconstructed with superior quality.

The second scene (Figure 9) contains examples for much more complicated transport paths. The scene is rendered in a set of environment maps. The figure demonstrates that we synthesize intermediate reflectance samples for all illumina-

tion directions in the hemisphere. The first column shows the result obtained from linearly blending between the input images, the images in the center column are synthesized based on the initial reconstruction R^{reco} with highlights added in. In the third column we present results rendered with our final reconstruction after regularization with image-based priors. The image-based priors remove noise and other artefacts that occur at shadow boundaries in the initial reconstruction due to imprecise shadow segmentation and labeling. The improved quality due to the priors is best visible under illumination with a single point light (first two rows). When illuminated with a complex environment map the noise in the initial reconstruction already averages out, and the expensive post-processing using image-based priors might be skipped, though the shadow boundaries are still slightly sharper in the third column. We provide renderings of rotating environment maps in the attached video.

We now analyze the differences between linear blending and our interpolation algorithm. The mirror sphere in the front reflects the environment and shows strong interreflections with the neighboring textured, slightly more glossy sphere. Extracting and warping the highlight maps, a sharp reflection of the environment is rendered with our approach (second and third column) while linear blending clearly produces artefacts revealing the original input samples even on the glossy sphere. With our method for shadow extraction, blending and reconstruction, we are able to correctly move and interpolate shadows avoiding the triple contours and banded shadows that are visible in all renderings using the linear blending. The shadow cast by the two spheres onto the postcard moves with the light source direction while the texture itself stays fixed.

The selection of different interpolation schemes derived from the warped shadow maps (Section 6.2) is powerful enough to even cope with non-trivial shadows as cast by the semi-transparent refracting spheres in the front right and in the back. They produce shadows tinted by the color of the spheres and furthermore exhibit an interesting intensity variation within the shadowed regions caused by caustics. Both the relatively diffuse caustic of the blue sphere in the front and the sharp caustic of the sphere behind the glass are faithfully reproduced with our method, moving smoothly. In contrast, they fade in and out in the linear blend.

7.1. Limitations

We observed some artefacts in the reconstruction of the shadow and caustics of the champagne glass. While the caustics have high spatial frequency, they are relatively dim. Our algorithm fails to detect them as highlights which need to be flowed, and therefore produces an incorrect interpolation. Improving on the segmentation algorithm, which is currently based on two manually selected thresholds f_{lit} and $f_{highlight}$, could help in this respect.

Another problem that is mostly noticeable at grazing angles is that illumination features, e.g. caustics or shadow

edges, might move too far in image space between two neighboring measured light samples to be handled correctly by our interpolation scheme. If there is no sufficient overlap in the shadow regions the applied level-set blending cannot produce the expected smooth motion. Capturing more samples in these regions would help.

The combined effect is visible for the champagne glass under point light source illumination where the complicated structures within in the shadow are not as expected, but they do not show the banding effects visible in the linear blend either. When illuminated by a complex environment map, the artifacts are barely visible.

8. Conclusion and Future Work

In this paper we have presented a novel and powerful interpolation framework for 4D reflectance fields. Our algorithm augments sparsely sampled reflectance fields by synthetic intermediate slices, simulating the displacement of highlights and shadows in the image plane as the direction of incident light changes. We have demonstrated that the augmented reflectance field can be used for creating realistic images both for extended light sources and for directional light.

The method interpolates highlights, cast and attached shadows in a plausible way, and is flexible enough to interpolate even non-trivial illumination effects such as shadows cast by semi-transparent or translucent objects. As we only use local pixel neighborhoods to determine whether a pixel is shadowed or not, but otherwise decouple its local appearance from the appearance of nearby pixels, we can precisely maintain texture information along shadow boundaries (such as on the postcard) without dragging it along as the shadow moves. Our method does not require any previous knowledge on the scene geometry and operates only on the reflectance field data. At the same time, the reduced number of samples makes the measurement significantly faster. We hope that these developments help to make reflectance fields applicable to a wide range of new applications in the future.

The output of our algorithm is currently a reflectance field of full resolution which consumes significant memory space. As the output is deterministically generated from the sparse input data, it is largely redundant. While it is not economical to re-run the algorithm for each rendered illumination situation, a first step towards successful compression could be to exploit that most pixels are computed as linear combinations of the neighboring input images, and store their location in a compressed mask. This could also allow for faster rendering times, as the integration of linearly changing reflectance values can be sped up.

Due to the demanding computation for a full reflectance field with up to 4000 intermediate images, we currently apply the image-based priors in a post processing step only. In the future we plan to integrate image-based priors in every step of the pipeline to avoid errors in the labeling and in the estimated flow.

9. Acknowledgements

We would like to thank Thomas Brox for providing an implementation of his flow algorithm, Ken Clarkson, who wrote the program `hull` which we used for the triangulation computation, and Carsten Stoll for his help with the video. This work has been partially funded by the Max Planck Center for Visual Computing and Communication (BMBF-FKZ01IMC01).

References

- [AF03] A.W. FITZGIBBON Y. WEXLER A. Z.: Image-based rendering using image-based priors. In *Proc. International Conference on Computer Vision (ICCV)* (2003), pp. 1176–1183. 3
- [BA96] BLACK M. J., ANANDAN P.: The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding, CVIU* 63, 1 (Jan. 1996), 75–104. 2
- [BBPWk04] BROX T., BRUHN A., PAPANBERG N., WEICKERT J.: High accuracy optical flow estimation based on a theory for warping. In *European Conference on Computer Vision (ECCV)* (Prague, Czech Republic, May 2004), Pajdla T., Matas J., (Eds.), vol. 3024 of *LNCS*, Springer, pp. 25–36. 2, 4
- [BT78] BARROW H., TENENBAUM J.: Recovering intrinsic scene characteristics from images. In *Computer Vision Systems*, Hanson A., Riseman E., (Eds.). Academic Press, 1978. 2
- [CL05] CHEN B., LENSCH H. P. A.: Light Source Interpolation for Sparsely Sampled Reflectance Fields. In *Proc. Vision, Modeling and Visualization* (2005), pp. 461–469. 3
- [Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of SIGGRAPH* (July 1998), pp. 189–198. 8
- [DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *Proceedings of ACM SIGGRAPH 2000* (July 2000), Computer Graphics Proceedings, Annual Conference Series, pp. 145–156. 2
- [FLBSar] FUCHS M., LENSCH H. P. A., BLANZ V., SEIDEL H.-P.: Adaptive sampling of reflectance fields. *ACM Transactions on Graphics* (to appear). 1
- [GTLL06] GARG G., TALVALA E.-V., LEVOY M., LENSCH H. P. A.: Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Rendering Techniques 2006 (Proc. of Eurographics Symposium on Rendering)* (June 2006), pp. 251–262. 2
- [HED05] HAWKINS T., EINARSSON P., DEBEVEC P.: A dual light stage. In *Rendering Techniques 2005 (Proc. of Eurographics Symposium on Rendering)* (June 2005), pp. 91–98. 2
- [LLSS03] LENSCH H. P. A., LANG J., SÁ A. M., SEIDEL H.-P.: Planned sampling of spatially varying brdfs. *Computer Graphics Forum (Proc. of Eurographics)* 22, 3 (Sept. 2003), 473–482. 2
- [MDA02] MASSELUS V., DUTRÉ P., ANRYS F.: The Free-form Light Stage. In *Proc. EuroGraphics Workshop on Rendering* (Pisa, Italy, June 26–28 2002), Debevec P. Gibson S., (Ed.), pp. 257–265. 2
- [MKL*04] MATSUSHITA Y., KANG S.-B., LIN S., SHUM H.-Y., TONG X.: Lighting and shadow interpolation using intrinsic lumigraphs. *Journal of Image and Graphics (IJIG)* 4, 4 (Oct. 2004), 585–604. 3
- [MPDW03] MASSELUS V., PEERS P., DUTRÉ P., WILLEMS Y. D.: Relighting with 4d incident light fields. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 22, 3 (2003), 613–620. 2
- [MPDW04] MASSELUS V., PEERS P., DUTRÉ P., WILLEMS Y. D.: Smooth reconstruction and compact representation of reflectance functions for image-based relighting. In *Rendering Techniques 2004 (Proc. of Eurographics Symposium on Rendering)* (June 2004), pp. 287–298. 2, 3
- [MPZ*02] MATUSIK W., PFISTER H., ZIEGLER R., NGAN A., MCMILLAN L.: Acquisition and Rendering of Transparent and Refractive Objects. In *Thirteenth Eurographics Workshop on Rendering* (June 2002), pp. 277–288. 2
- [PD05] PEERS P., DUTRÉ P.: Inferring reflectance functions from wavelet noise. In *Rendering Techniques 2005 (Proc. of Eurographics Symposium on Rendering)* (June 2005), pp. 173–182. 2
- [RB06] ROTH S., BLACK M. J.: Specular flow and the recovery of surface structure. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2006), vol. 2, pp. 1869–1876. 2
- [SCG*05] SEN P., CHEN B., GARG G., MARSCHNER S. R., HOROWITZ M., LEVOY M., LENSCH H. P. A.: Dual photography. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* 24, 3 (Aug. 2005), 745–755. 2
- [Whi00] WHITAKER R. T.: A level-set approach to image blending. *IEEE Transactions on Image Processing* 9 (Nov 2000), 1849–1861. 2, 4, 5
- [WSI04] WEXLER Y., SHECHTMAN E., IRANI M.: Space-Time Video Completion. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'04)* (2004), vol. 1, pp. 120–127. 3, 4, 6, 7
- [ZREB06] ZICKLER T., RAMAMOORTHY R., ENRIQUE S., BELHUMEUR P. N.: Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28, 8 (2006), 1287–1302. 2