Data-intensive Image based Relighting

Biswarup Choudhury* Sharat Chandran†

Vision, Graphics and Imaging Laboratory (ViGIL)
Department of Computer Science and Engineering
IIT Bombay

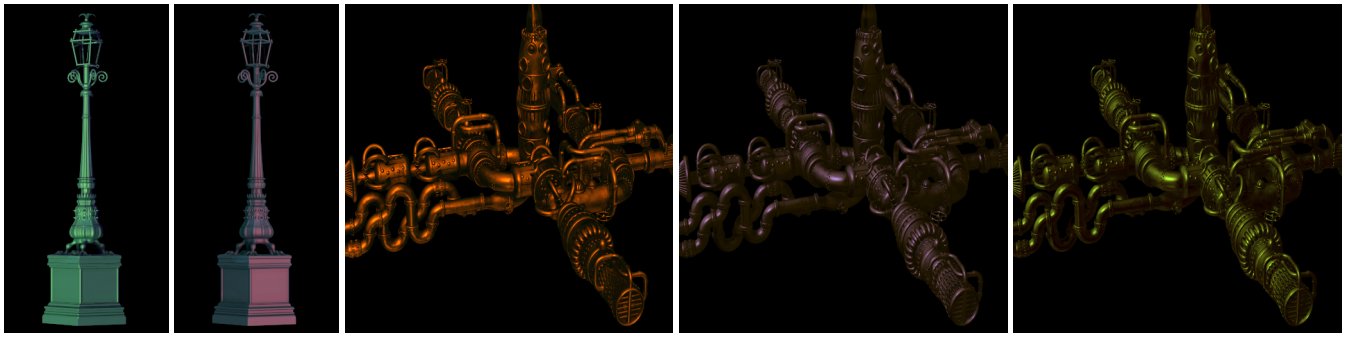

Figure 1: *LampPost and PipeSet relit using a light source with novel direction, color and intensity.*

Abstract

Image based Relighting(IBRL) has attracted a lot of interest in the computer graphics research, gaming, and virtual cinematography communities for its ability to relight objects or scenes, from novel illuminations captured in natural or synthetic environments. However, the advantages of an image-based framework conflicts with a drastic increase in the storage caused by the huge number of reference images pre-captured under various illumination conditions. To perform fast relighting, while maintaining the visual fidelity, one needs to preprocess this huge data into an appropriate model.

In this paper, we propose a novel and efficient two-stage relighting algorithm which creates a compact representation of the huge IBRL dataset and facilitates fast relighting. In the first stage, using Singular Value Decomposition, a set of eigen image bases and relighting coefficients are computed. In the second stage, and in contrast to prior methods, the correlation among the relighting coefficients is harnessed using Spherical Harmonics. The proposed method thus has lower memory and computational requirements. We demonstrate our results qualitatively and quantitatively with new generated image data.

Keywords: relighting, image-based, virtual/augmented reality

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality; I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Reflectance

*e-mail: biswarup@cse.iitb.ac.in
†e-mail: sharat@cse.iitb.ac.in

1 Introduction

The traditional way of synthesizing interesting imagery involves specifying all the objects in the world and their interactions, which can be termed the *source description*. An alternative way to describe the world is through the *appearance description*, also termed Image-based Modeling and Rendering(IBR). Unlike traditional rendering, IBR synthesizes realistic images from pre-recorded imagery without a complex and long rendering process.

A key area of interest in computer graphics has been illumination changes or relighting. Lighting design is one the most important decisions artists and designer have to take, both for our real and virtual world. Unfortunately, the ability to control illumination changes is inherently difficult with pre-acquired images. If this process of relighting can be made independent of the scene complexity, as in image-based relighting(IBRL), the artist is saved an enormous amount of time fine-tuning the illumination conditions to create the desired effect. For completeness, we give a brief description of the basic concepts of image-based relighting.

1.1 Image-based Relighting

Adelson [1991] proposed a seven dimensional *plenoptic function* for evaluating the low-level human vision models. The plenoptic function models a 3-dimensional environment by recording the light rays at every space location (V_x, V_y, V_z), from every possible direction (θ, ϕ), over any range of wavelengths (λ) and at any time (t), i.e.,

$$P = P(V_x, V_y, V_z, \theta, \phi, \lambda, t) \quad (1)$$

Although the definition of the plenoptic function is very general, it encapsulates all the dynamic effects, like motion, lighting changes etc. into this single parameter t. Wong et al. [Wong and Heng 2004] extends this plenoptic function to support relighting by extracting the illumination component (L) from the aggregate time parameter and explicitly specify it in the new formulation,

$$P = P(L, V_x, V_y, V_z, \theta, \phi, \lambda, t') \quad (2)$$

where, t' is the time parameter after extracting out the illumination component.

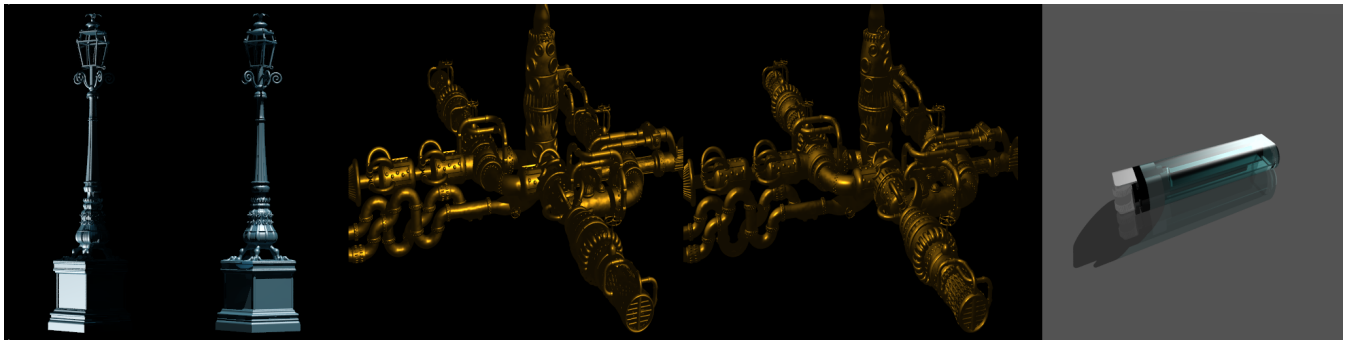

Figure 2: *Snapshots of the LampPost, PipeSet and Lighter datasets under different lighting directions, sampled on a spherical grid (Fig 8).*

Intuitively, the function tells us how the environment looks like when illuminated by a distant point light source. The authors call it the *plenoptic illumination function*. Since the newly introduced dimension L is a direction, sampling this function along this dimension is about capturing images under different light directions. Given a novel light vector, the desired image can be estimated by interpolating samples, hence relighting the scene. A similar, but differently parameterized formulation, *Reflectance function* [Debevec et al. 2000] is a 8-dimensional function which takes measurement of how materials transform incident illumination into radiant illumination. Relighting is performed by multiplying the coefficients of the reflectance function and the incident illumination.

The approach chosen by many IBRL techniques involves pre-rendering (synthetic) or pre-acquisition (real) of a collection of images in which the lighting direction is systematically varied. If the density of illumination is dense enough, then due to linearity of scene radiance, images of the scene under complex illumination can be computed simply by superposition of single light source images. Although now relighting is tractable, the collection of images is typically too large both to store in memory and to synthesize novel images in real-time (Debevec et al. [Debevec et al. 2000] use 2000 images and Koudelka et al. [Koudelka et al. 2001] use more than 4000 images). If too few images are used, the quality of reconstruction is compromised (blurring of specularities and inaccurate shadows). The fundamental fact of IBRL data being closely related to the surface reflectance makes it necessary to correctly model the data into a compact and efficient representation. We discuss in the next section ways in which researchers have coped with this issue.

1.2 Related Work

Compression: Several compression techniques have been proposed to remove the data redundancy in image-based data. In [Debevec et al. 2000], images of each pixel's reflectance function are stored in the JPEG format and further processed in the compressed domain. Lin et al. [Lin et al. 2002], on the other hand uses a 2D DCT to compress the images of a pixel's radiance values. Vector quantization, entropy coding, and wavelet transform are some of the widely used image compression techniques. However, these techniques either compress the relighting data by small factor or (their overuse) introduces artifacts.

Sampling: As a novel image is synthesized from a model built from reference images, the quality of reconstruction depends on the sampling density of reference images. Lin et al. [Lin et al. 2002] derived a theoretical geometry-independent sampling bound for IBRL based on radiometric tolerance. This satisfies the sampling needs only of certain scenes; further, even with the guidance of a sampling bound, the resource requirements are huge.

Global dimensionality reduction: Hallinan [Hallinan 1994] and Ho et al. [Ho et al. 2005] used Principal Component Analysis(PCA) to model for the variation of images due to illumination changes. Nimeroff et al. [Nimeroff et al. 1994] showed that linear bases (steerable functions) could be used for relighting of scenes under complex, but diffuse illumination. Several variants (refer [Choudhury and Chandran 2006]) of representing variation of images due to illumination has been analyzed and proposed in the past. The desired images can be synthesized by linearly combining "eigen-images" given a set of eigen-coefficients. Unfortunately, eigen-coefficients cannot be related to a lighting direction by a simple mathematical equation.

Wong et al. [Wong et al. 1997a; Wong and Leung 2003] and Sloan et al. [Sloan et al. 2002] use spherical harmonics(SH) to model the data. An interesting note about spherical harmonics is that the corresponding harmonic basis functions are analytical functions and hence they do not need to be pre-computed and stored. On a different note, Leung et al. [Leung et al. 2006] uses spherical radial basis functions(SRBF) to model the IBRL data. Since SRBF is prone to quantization noise, a complex constrained optimization problem is solved for iteratively, to generate the SBRF weights. Though the computation of SRBF functions are faster, the quality of results achieved is lower than that using SH.

Regardless of the type of linear bases used, greater accuracy, in the realm of low frequency lighting (sharp shadows and specularities), in the renderings requires many bases. However, as the number of bases increases so do the storage requirements and rendering time.

Local dimensionality reduction: When a complete image is considered, complex illumination effects such as cast shadows, specularities and caustics cannot be represented by a handful of linear basis functions. However, if these illumination effects are analyzed locally, i.e., at small regions within the image, it is possible to model the variation within these regions ([Matusik et al. 2002; Nayar et al. 2004; Nishino et al. 2001; Wood et al. 2000]). For a region of a scene with simple illumination effects, lesser number of bases would be needed than a region with complicated reflectances. Since we would like to efficiently relight objects of all possible reflectances, therefore we chose this methodology for our relighting algorithm. For the interested readers, refer [Choudhury and Chandran 2006] for a detailed survey of IBRL techniques.

The rest of the paper is organized as follows, Contributions of our work are stated in Section 1.3. In Section 2, we describe our novel relighting algorithm; specifically, Section 2.1 and Section 2.2 discusses the two stages of our algorithm respectively, and Section 2.3 presents our final relighting (reconstruction) algorithm. Section 3 discusses the results generated and provides an overall evaluation of our algorithm. Finally, we provide our conclusions in Section 4.

1.3 Contributions

In this paper, we tackle the complex problem of relighting a scene with novel illumination. We provide an efficient and novel image based relighting technique which exploits the above observations (of local dimensionality reduction) and efficiently produces fast realistic renderings. Specifically,

- We propose a two-stage algorithm for simultaneously exploiting the correlations in the lighting domain and spatial domain inherent in IBRL datasets. The output of the above algorithm, image bases and relighting coefficients, are stored and used for fast relighting.
- We provide three synthetic IBRL datasets, *LampPost*, *PipeSet* and *Lighter* (Figure 2) which enables validation of our technique. The datasets are available for viewing and download at `www.cse.iitb.ac.in/biswarup/web/data/`

2 The Algorithm

The input is a set of reference images with the same viewpoint, but under different lighting directions sampled uniformly on a sphere (Halton sampling [Wong et al. 1997b]) as illustrated in Figure 8. Figure 2 shows example images of our dataset (*LampPost*, *PipeSet* and *Lighter*), each illuminated by a single distant point light source. Notice the specularities, highlights, and shadows in the renderings.

Our algorithm is composed of two components. First, exploiting the correlation among pixels of an image, we compute a set of image bases and their corresponding relighting coefficients. Second, exploiting the coherence among the computed relighting coefficients we produce a reduced set of (spherical harmonic) relighting coefficients. Figure 3 diagrammatically illustrates our algorithm. Given a novel lighting configuration, using the "lightweight" pre-calculated image bases and spherical harmonic relighting coefficients, we efficiently reconstruct (relight) the scene.

2.1 Stage 1: Singular Value Decomposition

We first reorganize the input data to suit the needs of subsequent processing. We have as input, a collection of n images of a scene under varying lighting directions. Each is an image illuminated by a single distant point source. We compute a mean image over all the reference images and subtract it from each of them to generate a set of images I_i. The mean extraction offers a more accurate computational base for further processing. Next, we divide all images into m square blocks, each block containing p pixels (and then linearize each of the m square blocks into a data vector). We now, create a $n \times p$ matrix I^j as a collection of image blocks in which the k^{th} row corresponds to the j^{th} block of image I_k. After computing a low-dimensional approximation of each such block matrix I^j in the scene using Singular Value Decomposition(SVD), we find a rank b approximation to I^j ($b \ll n$) as

$$I^j \approx U^j W^j V^j \tag{3}$$

where U^j is a $n \times b$ matrix, W^j is a $b \times b$ diagonal matrix and V^j is a $b \times p$ matrix.

We now split the singular values into two halves (taking square roots) and multiply them to U^j and V^j, thereby evenly distributing the energy of singular values. We are therefore able to rewrite Eqn. 3 as

$$I^j \approx R^j E^j \tag{4}$$

where R^j is a $n \times b$ matrix and E^j is a $b \times p$ matrix.

Let the matrix $\mathbf{R}$ ($n \times m \times b$) be formed by stacking all of the m R^j block matrices. We term $\mathbf{R}$ as the SVD relighting coefficient matrix. Likewise, let the collection of eigen bases for all blocks be denoted by $\mathbf{E}$ ($b \times m \times p$). We term $\mathbf{E}$ as the eigen (image) bases. The choice of the above terminology (for $\mathbf{R}$ and $\mathbf{E}$) is to provide a physical interpretation of the SVD decompostion. Since the dimension of an input reference image and an element of (the stacked up) $\mathbf{E}$ is the same (Figure 3), $b \times m \times p$, it is appropriate to term $\mathbf{E}$ as eigen (image) bases. Similarly, since $\mathbf{R}$ is a $n \times m \times b$ matrix, its elements are indexed by the n lighting vectors and hence is appropriately termed the relighting coefficients.

The creation of relighting coefficients $\mathbf{R}$ and eigen bases $\mathbf{E}$ is illustrated diagrammatically in Figure 3. The collection of submatrices within $\mathbf{R}$ and $\mathbf{E}$ contain all the information needed to approximate the collection of images corresponding to all the n lighting directions.

Discussion: The number of eigen bases (b) depends on the surface reflectance and geometry of the scene. For determining the number of eigen bases, we measured the peak signal-to-noise ratio(PSNR) of the reconstructed images (by a multiplication of $\mathbf{R}$ and $\mathbf{E}$). From experimentation, we concluded that for correct relighting, the value of b should be such that the variance of the decomposition is 70%–90%.

Before the blockwise SVD procedure was applied, we had nmp measurements and after the SVD, we have bmp elements in $\mathbf{E}$ (considering only the eigen bases), thus providing a compression ratio of b/n; in general this is a significant saving since, for blockwise SVD, we expect $b \ll n$ with little compromise in the fidelity to the original data. In the specific case of *LampPost*, the number of images $n = 1000$, the number of blocks $m = 128$, the number of pixels per block $p = 1024$ and the number of bases b per block ranges from 10–300.

By using blocks, we have introduced the additional burden of keeping track of $\mathbf{R}$, consisting of nmb elements. In general, the storage requirements of $\mathbf{R}$ are considerable. For example, in the case of *LampPost*, 302MB was needed to store the eigen bases and relighting coefficients. The extra burden of keeping track of the relighting coefficient matrix $\mathbf{R}$ is tackled by exploiting the fact that there is inter-block coherence among the m different image blocks. Therefore, it is likely that the relighting coefficients for each of the blocks are not linearly independent.

2.2 Stage 2: Spherical Harmonics

The relighting coefficients in $\mathbf{R}$ are indexed by the light vectors (θ, ϕ), or rephrasing, $\mathbf{R}$ is a spherical function. Hence Spherical Harmonics can be used to harness the above mentioned inter-block correlation.

Spherical Harmonics: We transform the relighting coefficient matrix $\mathbf{R}$ using spherical harmonics and the resultant spherical harmonic coefficients are zonally sampled and quantized. Given the standard parametrization of points on the surface of a unit sphere into spherical coordinates $(\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$, the spherical harmonics function is represented as:

$$Y_{l,m}(\theta,\phi) \quad = \quad \begin{cases} P_{l,m}Q_{l,m}(\cos\theta)\cos(m\phi) & \text{if m>0;} \\ P_{l,m}Q_{l,-m}(\cos\theta)\sin(|m|\phi) & \text{if m<0;} \\ P_{l,0}Q_{l,0}(\cos\theta)/\sqrt{2} & \text{if m=0;} \end{cases}$$

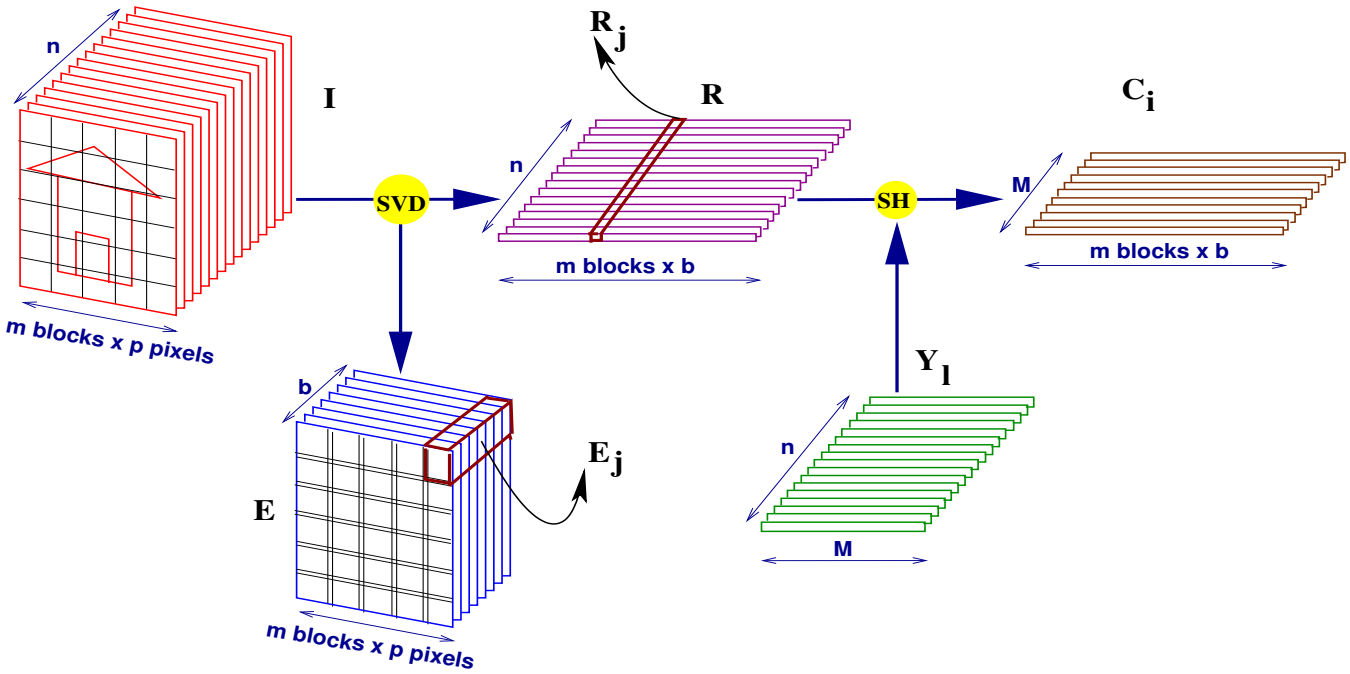

Figure 3: *Algorithm Flowchart: In the first stage (SVD), the input reference images are decomposed to eigen image bases* $\mathbf{E}$ *and relighting coefficients* $\mathbf{R}$ *(Section 2.1). The coherence among relighting coefficients* $\mathbf{R}$ *are further harnessed using Spherical Harmonics (Section 2.2).*

$$P_{l,m} \quad = \quad \sqrt{\frac{(2l+1)}{(2\pi)}\frac{(l-|m|)!}{(l+|m|)!}}$$

and

$$Q_{l,m}(x) = \begin{cases} (1-2m)\sqrt{1-x^2}Q_{m-1,m-1}(x) & \text{if l=m;} \\ (2m+1)xQ_{m,m}(x) & \text{if l=m+1;} \\ \frac{2l-1}{l-m}xQ_{l-1,m}(x) - \frac{l+m-1}{l-m}Q_{l-2,m}(x) & \text{otherwise;} \end{cases}$$

and $Q_{0,0}(x) = 1$.

$Q_{l,m}$ are the associated Legendre polynomials, $P_{l,m}$ is a scaling factor and l is the degree of the spherical harmonic representation. Please note that m, in the context of SH, is in no way related to the m blocks in any image I_i (Section 2.1). The spherical function is projected into spherical harmonic coefficients by integrating the product of the former and the spherical harmonic function $Y_{l,m}$. This integral is evaluated using Monte Carlo integration.

$$\mathbf{C}_{l,m} \quad = \quad \int_0^{2\pi}\int_0^{\pi} \mathbf{R}(\theta,\phi)Y_{l,m}(\theta,\phi)\sin\theta d\theta d\phi \qquad (5)$$

or alternatively [Green 2003],

$$\mathbf{C}_k \quad = \quad \frac{4\pi}{n}\sum_{i=1}^{n}\mathbf{R}(x_i)Y_k(x_i) \qquad (6)$$

where n is the number of lighting directions (as before) in the illumination space (θ,ϕ) and x_i is the flattened array of these illumination directions. $\mathbf{R}(\theta,\phi)$ (or $\mathbf{R_k}$) are the sampled SVD relighting coefficients, coefficients $\mathbf{C}_{l,m}$ (or $\mathbf{C_k}$) are the spherical harmonic coefficients to be zonally sampled, quantized and stored for relighting. To reconstruct the approximated relighting coefficients $\mathbf{R}^*(\theta,\phi)$, we reverse the process and sum scaled copies of the corresponding spherical harmonic functions.

$$\mathbf{R}^*(\theta,\phi) \quad = \quad \sum_{l=0}^{l_{\max}}\sum_{m=-l}^{l}\mathbf{C}_{l,m}Y_{l,m}(\theta,\phi) \qquad (7)$$

where $(l_{\max}+1)^2$ is the total number of SH coefficients stored; for Eqn. 6,

$$\mathbf{R}^*(x_i) \quad = \quad \sum_{k=1}^{M}\mathbf{C}_k Y_k(x_i) \qquad (8)$$

where $M = (l_{\max}+1)^2$, is the total number of SH coefficients stored.

Discussion: Increasing the number of SH relighting coefficients increases the fidelity, but also with increasing computation and storage requirements. The low-order spherical harmonics are responsible for the diffuse component of illumination. On the other hand, the higher-order harmonics are responsible for specularities, highlights and shadows. The optimal number of SH coefficients therefore depends on the image content.

To objectively evaluate the quality of spherical harmonic compression, we measure the peak signal-to-noise ratio(PSNR) against the number of coefficients used for encoding. We observed that in most cases, a degree (l_{max}, Eqn. 7) of 10–25 was enough to maintain a fair balance between visual fidelity, storage and computation.

Before applying SH on $\mathbf{R}$, we had nmb measurements and after SH, we have Mmb (Eqn. 8) elements in $\mathbf{C}_{l,m}$ (or $\mathbf{C}_k$). We get a theoretical compression gain of n/M. In the case of *LampPost*, the number of images $n = 1000$ and M was in the range of 100–600, depending on the illumination effects.

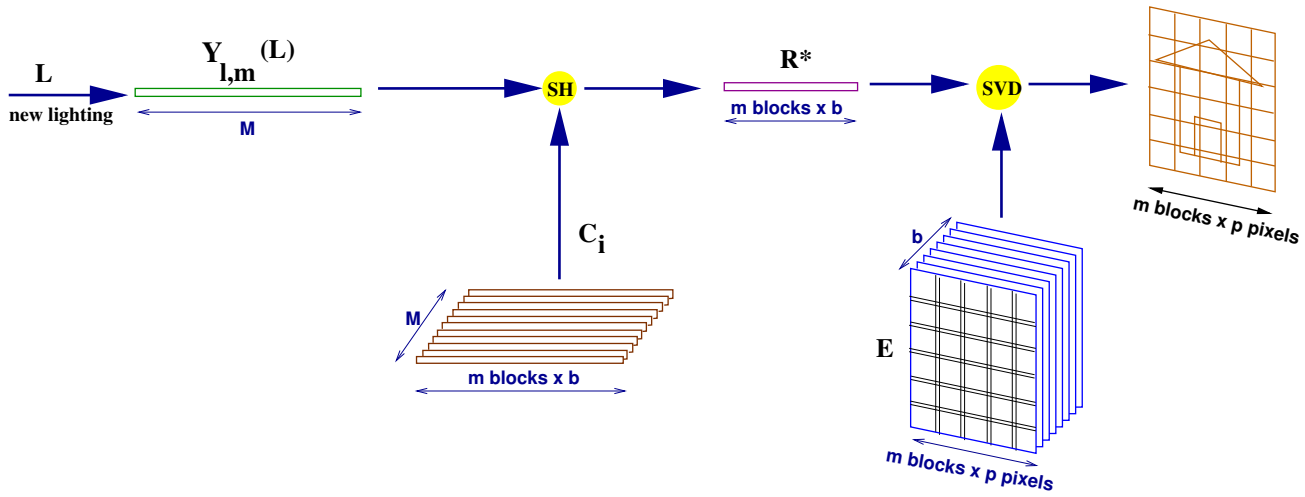

Figure 4: *Illustration of a typical reconstruction under a novel lighting L. Given L, the SH basis functions are calculated, multiplied with the stored SH relighitng coefficients, to generate the appropriate relighting coefficient* $\mathbf{R}$*. The final relit image is computed by multiplying* $\mathbf{R}$ *with the eigen image bases.*

2.3 Fast Relighting

All the data needed to finally render the image under a new specified illumination are stored in the matrices $\mathbf{E}$ and $\mathbf{C}_i$. These matrices are compact, have low storage requirements and facilitate fast relighting of the scene. Here, we detail the steps needed to render an novel image under a new specified (directional) light source $L(\theta', \phi')$.

To render an image of the scene as if it were illuminated by a new light source $L(\theta', \phi')$, we first calculate the spherical harmonic functions $Y_{l,m}(\theta', \phi')$ and then compute the product $C_{l,m}Y_{l,m}(\theta', \phi')$, the reconstructed relighting coefficients ($\mathbf{R}^*$). We then compute the product $\mathbf{R}^*\mathbf{E}$, which essentially is the novel image of the scene under the specified lighting (In Figure 4, we diagrammatically illustrate this process). In the case of using a high-dynamic range environment map as the new lighting source, each pixel (or a set of them) corresponds to one of the n lighting directions in the original collection of images.

Our algorithm can also be applied for relighting in the case of a positional light source, if the physical coordinate of the surface point associated with every pixel is available. Given the position of the positional light source, the corresponding lighting vector for every pixel can be calcuated [Wong and Heng 2004]. These lighting vectors are then used to compute spherical basis functions for every pixel, which further facilitates the computation of the corresponding relighting coefficients $\mathbf{R}$. Relighting can then be performed as specified above.

Color Images – In the case of color images, there is an image for every color channel and so, the number of images specified by n becomes $3n$. We perform both stages (SVD and SH) of the relighting process on each color channel separately. As there is a lot of redundancy in the color channels from image to image, we keep the number of eigen bases b and the number of SH relighting coefficients M fixed.

3 Implementation and Results

The two-stage relighting algorithm was implemented on a Intel 2.4 GHz processor with 1Gb of RAM. We generated two new IBRL datasets, *LampPost* and *PipeSet* for purposes of our experimentation. Both the datasets were generated by uniformly sampling a sphere into 1000 and 250 sample points respectively (see Figure 8) using Halton points [Wong et al. 1997b], which have been proved [Cui and Freeden 1997] to generate evenly distributed and stratified samples on the surface of a sphere. The rendering of the datasets was performed with radiosity computations using POV-Ray. The resolution of an image in the datasets *LampPost*, *PipeSet* is 512×256 and 512×512 respectively.

For purposes of comparison, we chose two commonly used standard image based relighting algorithms,

1. Two-stage SVD(2S-SVD) [Nayar et al. 2004]: The first stage of this technique is similar to the first stage (Section 2.1) of our algorithm. In its second stage, SVD is applied to capture the coherence among the relighting coefficients.

2. Illumination-adjustable Images(IAI) [Wong and Leung 2003]: In this technique, every pixel, illuminated with light directions sampled on a sphere, is considered a spherical function and spherical harmonics is used to model each of them.

We chose three new lighting configurations and rendered both our models, *LampPost* and *PipeSet* under the same using POV-Ray. Under the same three lighting configurations, we performed experiments using our and both the above mentioned algorithms. In section 3.1 and section 3.2, we detail the results obtained from these experiments. Figure 5 and Figure 6 illustrates some of the relit images corresponding to the experiments performed.

3.1 Our Algorithm vs. Two-stage SVD

All reference images were divided into blocks of size 32×32 before processing. For the purpose of evaluation and comparison, we implemented the Two-stage SVD. In Table 1, we observe that in the case of *LampPost* and *PipeSet*, our relighting technique has much lesser storage requirements(Size). We also perform much faster relighting as compared to Two-stage SVD. Note that the total relighting time scales up (for Two-stage SVD) when we consider multiple light sources for relighting. The time taken to pre-process(Pre-P) the input data is more in our case but it is a one-time offline process, hence is independent of the final relighting time.

	Our Algorithm			Two-Stage SVD			IAI		
	Size(Mb)	Pre-P.(sec)	Relight(sec)	Size	Pre-P.	Relight	Size	Pre-P.	Relight
LampPost 1	130	4206.3	7.2	477	3736.3	29.0	513	34069	28.38
LampPost 2	419	5883.3	15.8	477	3871.4	28.8	513	55873	38.52
LampPost 3	144	2448.5	5.3	477	2195.4	19.4	513	56850	40.52
PipeSet 1	127	649.9	17.4	336	630.5	25.2	347	4566	42.31
PipeSet 2	142	745.5	18.5	336	645.4	25.86	661	11328	60.9
PipeSet 3	129	445.9	10.6	336	466.2	18.0	402	3546	31.74

Table 1: *Detailed results of the experiments performed using* **Our Algorithm**, **Two-stage SVD** *and* **IAI**. *Pre-P.(sec) is the total time taken (in seconds) by each technique to pre-process the input images and create a corresponding appropriate model. Size(Mb), in each case, represents the size of the computed model from the images (after the pre-processing step). This (model) is kept in memory for relighting purposes. Relight(sec) is the time taken (in seconds) to perform final relighting of the scene under a novel illumination.*

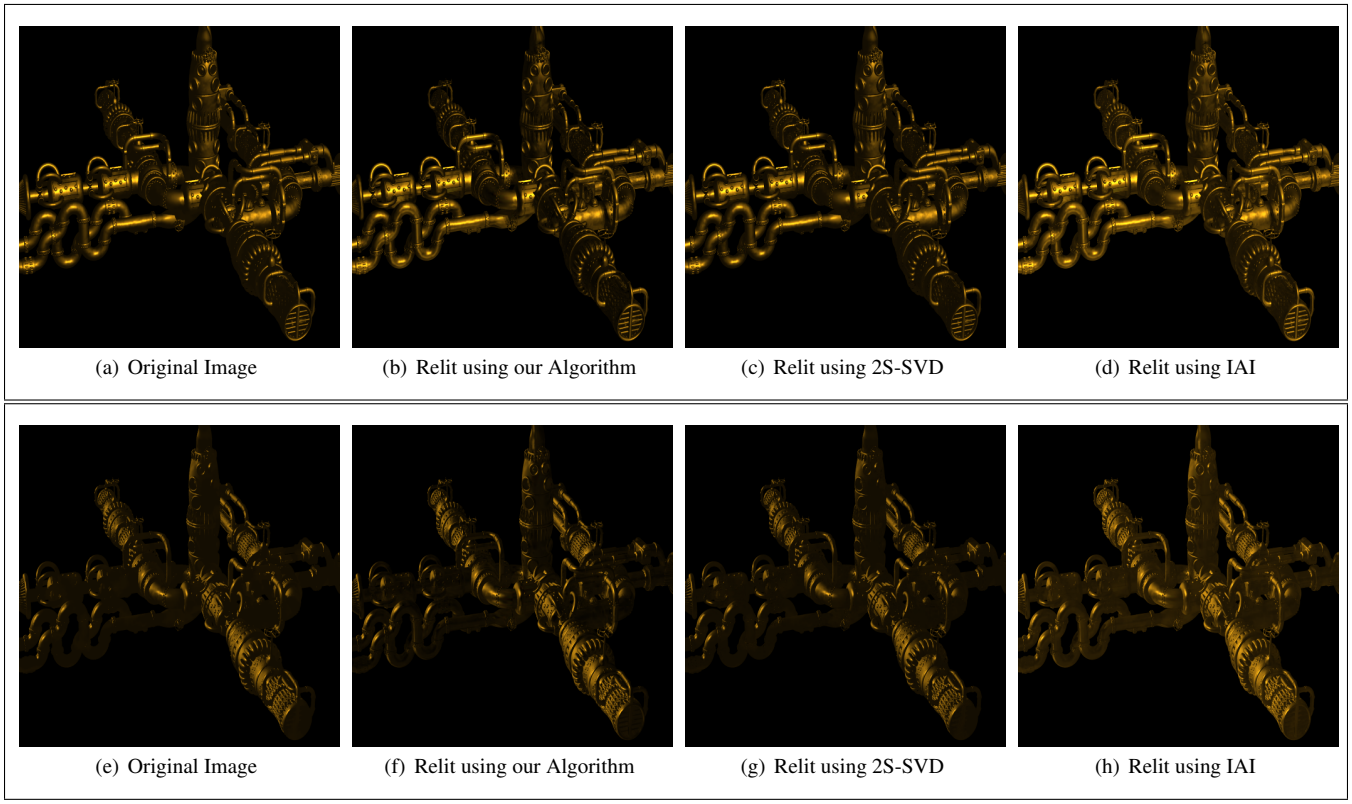

(a) Original Image (b) Relit using our Algorithm (c) Relit using 2S-SVD (d) Relit using IAI

(e) Original Image (f) Relit using our Algorithm (g) Relit using 2S-SVD (h) Relit using IAI

Figure 5: *Relighting PipeSet under two different novel illuminations. Figure 5(a),5(b),5(c),5(d) is a set under one illumination and, Figure 5(e),5(f),5(g),5(h) is the other set. Our algorithm produces an image which has almost no perceivable difference with the original image. Note the specular highlights, glossiness and shadows. For quantitative details of these results, see respective entries for Pipset 1 and Pipset 3 in Table 1.*

For the same number of bases, SVD has a slightly higher PSNR as compared to Spherical Harmonics. However, the tradeoff is in terms of, firstly, large storage requirements (in this case, two basis sets corresponding to the two SVD and a corresponding set of relighting coefficients) as compared to our algorithm (one basis set and a corresponding set of relighting coefficients), and secondly, slower final relighting.

3.2 Our Algorithm vs. IAI

We implemented [Wong and Leung 2003] (sans the discrete wavelet transform stage, to avoid slowing up relighting caused by the additional computation required during the decoding process) and used it for comparison. We observe that our algorithm clearly out performs (storage, final relighting and pre-processing time) IAI (see Table 1). IAI models every pixel (under input lighting conditions) as a spherical function using spherical harmonics. In other words, only the intra-pixel correlations are modeled. In addition, since spherical harmonics provides a suboptimal representation of the data, IAI needs to compute and store a huge number of relighting coefficients to achieve even comparable results in terms of the illumination effects and PSNR. On the other hand, our algorithm efficiently model the inter-pixel correlations in the first stage using SVD, and then model the intra-pixel coherence (among relighting coefficients) using spherical harmonics. Therefore, we need to store (comparitively) lesser number of bases images and relighting coefficients. This facilitates lesser storage requirements and enables faster relighting.

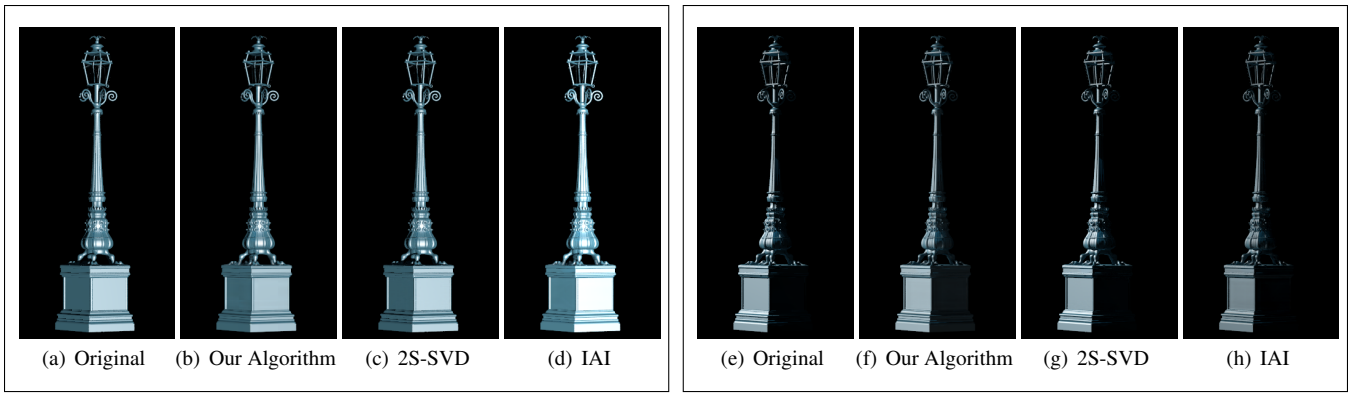

(a) Original (b) Our Algorithm (c) 2S-SVD (d) IAI (e) Original (f) Our Algorithm (g) 2S-SVD (h) IAI

Figure 6: *Relit LampPost under two different illuminations. Figure 6(a),6(b),6(c),6(d) are images under one lighting direction and, Figure 6(e),6(f),6(g),6(h) is the other set of images under a different lighting. Shadows, the desired sharpness and various illumination effects are faithfully reproduced. For quantitative details, see respective entries for* Lampost 1 *and* Lampost 2 *in Table 1.*

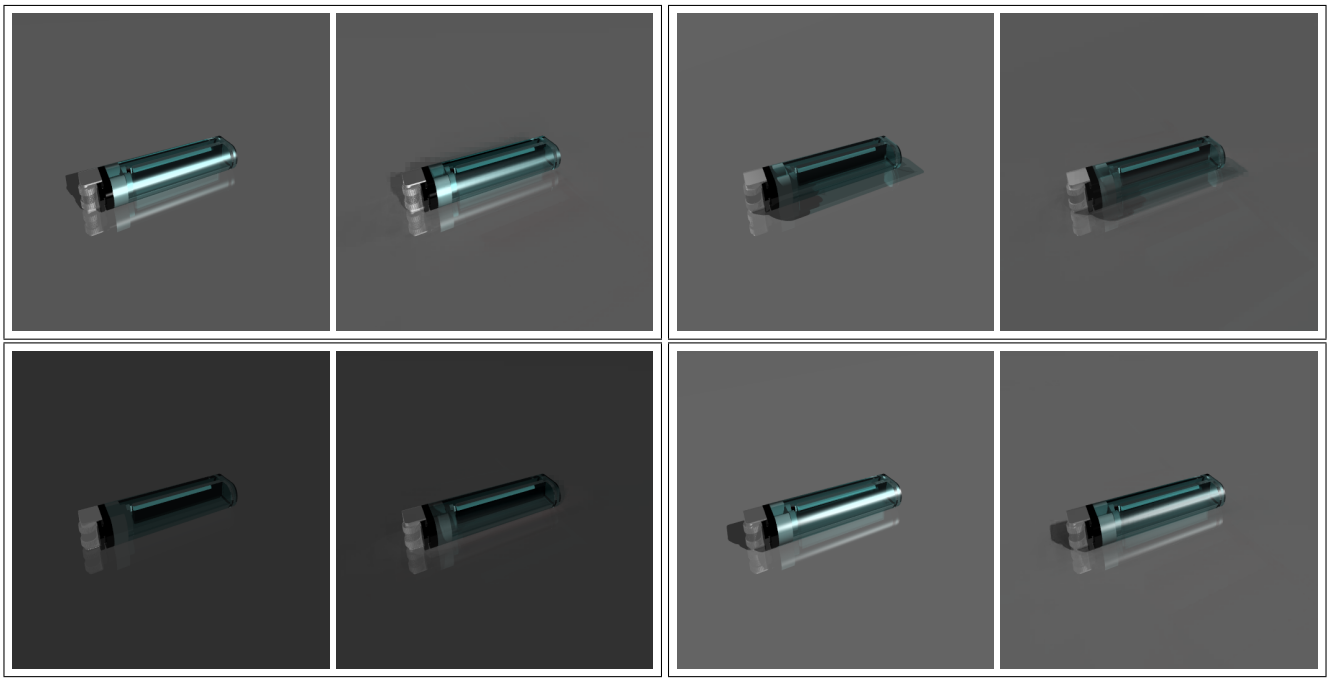

Figure 7: *Relit Lighter under four different illuminations. Each set consists of the original (left) and the relit image (right) of the Lighter under a different illumination direction. All relit images have PSNR>30. For details, refer Section 3.3.*

3.3 Discussion

For purposes of validation, we also generated and performed our experiments on an IBRL dataset of a translucent and refractive object *Lighter*. We sampled the *Lighter* under 1000 lighting directions and performed relighting on it under novel lighting directions (Figure 7). We do observe some blockiness in the shadows produced, though the other illumination effects are reproduced faithfully. Due to the complex nature of the *Lighter* and its illumination effects, we can tackle artifacts (and all effects) by increasing the input sampling density and the number of basis functions (both, for SVD and SH). To create correct and realistic rendering for all objects producing complex illumination effects, such as caustics, refraction, soft shadows, we sample the object densely and then process the data appropriately.

Our algorithm is amenable to a GPU implementation. Each block could be considered as an independent entity since all processing (both stages of our algorithm, SVD and SH and relighting) is performed on each block independently. Computing spherical basis functions takes some time, especially in the context of positional light sources where distinct spherical basis functions have to be computed for each pixel. However SH basis function evaluation can be really sped up by using hardware texture lookup (cube-maps) of a graphics processing units(GPU). Using these cubemaps and the inherent parallelism of the "state-of-the-art" GPUs, we could process all blocks simultaneously and thereby perform relighting in realtime.

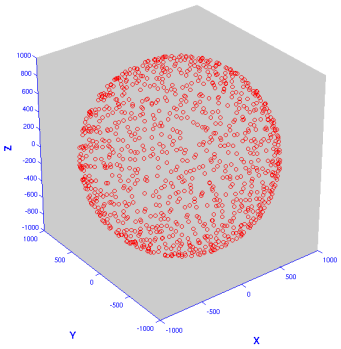

Figure 8: *Sample points on the surface of a sphere denotes point light sources used for illuminating the object/scene.*

4 Conclusion

In this paper, we propose a novel two-stage IBRL technique, which tackles the traditional problem of huge storage and computational resource requirements. We apply SVD to capture the inter-pixel correlations, producing a set of eigen bases images and corresponding relighting coefficients in Stage 1, and in Stage 2, we further model the intra-pixel correlations among the relighting coefficients using Spherical Harmonics and reduce them to a compact set of SH relighting coefficients. Fast relighting can then be performed with single/multiple light sources. Three new IBRL datasets, *LampPost*, *Pipset* and *Lighter* for the purpose of experimentation have been generated and experimental results validate our technique. The datasets are available at `http://www.cse.iitb.ac.in/biswarup/web/data/`.

Acknowledgements

We would like to thank the POV-Ray team for making available their free renderer and also, Micha Riser for maintaining the repository of free POV-Ray models at `http://objects.povworld.org/`. Special thanks to Deepali Singla for all the help extended during implementations and discussions. We also thank the members of ViGIL, IIT Bombay for their support. This research was supported by the fellowship fund endowed by Infosys Technologies Limited, India.

References

ADELSON, E. H., AND BERGEN, J. R. 1991. The plenoptic function and the elements of early vision. *Computational Models of Visual Processing*, 3–20.

CHABERT, C.-F., EINARSSON, P., JONES, A., LAMOND, B., MA, W.-C., SYLWAN, S., HAWKINS, T., AND DEBEVEC, P. 2006. Relighting human locomotion with flowed reflectance fields. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Sketches*, 76.

CHOUDHURY, B., AND CHANDRAN, S. 2006. A survey of image-based relighting techniques. In *GRAPP '06*, 176–183.

CUI, J., AND FREEDEN, W. 1997. Equidistribution on the sphere. *SIAM Journal of Science Computing 18*, 2, 595–609.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *SIGGRAPH '00*, 145–156.

GREEN, R. 2003. Spherical harmonic lighting: The gritty details. Tech. rep., Sony Computer Entertainment America.

HALLINAN, P. 1994. A low-dimensional representation of human faces for arbitrary lighting conditions. In *CVPR '94*, 995–999.

HO, P., WONG, T., AND LEUNG, C. 2005. Compressing the illumination-adjustable images with principal component analysis. *IEEE Transactions on Circuits and Systems for Video Technology 15*, 3, 355–364.

KOUDELKA, M. L., BELHUMEUR, P. N., MAGDA, S., AND KRIEGMAN, D. J. 2001. Image-based modeling and rendering of surfaces with arbitrary brdfs. In *CVPR '01*, 568–575.

LEUNG, C.-S., WONG, T.-T., LAM, P.-M., AND CHOY, K.-H. 2006. An RBF-based compression method for image-based relighting. *IEEE Transactions on Image Processing 15*, 4, 1031–1041.

LIN, Z., WONG, T.-T., AND SHUM, H.-Y. 2002. Relighting with the reflected irradiance field: Representation, sampling and reconstruction. *International Journal of Computer Vision 49*, 2-3, 229–246.

MATUSIK, W., PFISTER, H., NGAN, A., BEARDSLEY, P., ZIEGLER, R., AND MCMILLAN, L. 2002. Image-based 3d photography using opacity hulls. In *SIGGRAPH '02*, 427–437.

NAYAR, S. K., BELHUMEUR, P. N., AND BOULT, T. E. 2004. Lighting sensitive display. *ACM Transaction on Graphics 23*, 4, 963–979.

NIMEROFF, J. S., SIMONCELLI, E., AND DORSEY, J. 1994. Efficient Re-rendering of Naturally Illuminated Environments. In *Proceedings of the Eurographics Workshop on Rendering '94*, 359–373.

NISHINO, K., SATO, Y., AND IKEUCHI, K. 2001. Eigen-texture method: Appearance compression and synthesis based on a 3d model. *IEEE Transanction on Pattern Analysis and Machine Intelligence 23*, 11, 1257–1265.

SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *SIGGRAPH '02*, 527–536.

WONG, T.-T., AND HENG, P.-A. 2004. Image-based relighting: representation and compression. *Integrated image and graphics technologies*, 161–180.

WONG, T., AND LEUNG, C. 2003. Compression of illumination-adjustable images. *IEEE Transactions on Circuits and Systems for Video Technology 13*, 11, 1107–1118.

WONG, T.-T., HENG, P.-A., OR, S.-H., AND NG, W.-Y. 1997. Image-based rendering with controllable illumination. In *Proceedings of the Eurographics Workshop on Rendering Techniques '97*, 13–22.

WONG, T.-T., LUK, W.-S., AND HENG, P.-A. 1997. Sampling with hammersley and halton points. *Journal of Graphics Tools 2*, 2, 9–24.

WOOD, D. N., AZUMA, D. I., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D. H., AND STUETZLE, W. 2000. Surface light fields for 3d photography. In *SIGGRAPH '00*, 287–296.