

The Free-form Light Stage

Vincent Masselus, Philip Dutré, Frederik Anrys[†]

Department of Computer Science, Katholieke Universiteit Leuven, Belgium

Abstract

We present the Free-form Light Stage, a system that captures the reflectance field of an object using a free-moving, hand-held light source. By photographing the object under different illumination conditions, we are able to render the object under any lighting condition, using a linear combination of basis images. During the data acquisition, the light source is moved freely around the object and hence, for each picture, the illuminant direction is unknown. This direction is estimated automatically from the images. Although the reflectance field is sampled non-uniformly, appropriate weighting coefficients are calculated. Using this system, we are able to relight objects in a convincing and realistic way.

Categories and Subject Descriptors (according to ACM CCS): I.2.10 [Artificial Intelligence]: Vision and scene Understanding I.3.7 [Computer Graphics]: Three dimensional Graphics and Realism I.4.1 [Image processing and Computer Vision]: Digitization and Image Capture

1. Introduction

Relighting of objects or scenes has become an important research topic over recent years. Applications range from real-time global illumination and lighting design to mixed reality where real and virtual objects are combined with consistent illumination.

In this paper, we present the Free-form Light Stage, a system that captures the reflectance field of an object using a free-moving, hand-held light source. The reflectance field of an object is the excitant radiant light field of that object under any incident light field. The main idea in our object relighting technique is to create a linear combination of basis images. Each basis image is a photograph of the object, from a fixed viewpoint, with the light source positioned freely over the object. The light source direction is extracted from each photograph and used together with an incident light field to calculate the scalars in the linear combination. To record the photographs, a very simple system setup is used. We use a digital camera, a hand-held light source, four white diffuse spheres that will be used for light source extraction, and the object to be relit.

The structure of this paper is as follows: In section 2, we

overview some related work and position our technique in the broad field of relighting. We continue with an overview of the method in section 3. In section 4, the setup of the camera, positioning of the object and the diffuse spheres is surveyed. We overview the calibration of camera and light source as well, and the data acquisition process is considered in more detail.

For each photograph in the acquired set, we estimate the light source direction by analyzing the shading pattern of the diffuse spheres, visible in the photographs. Our light source direction estimation technique is explained in section 5.

The weights for the linear combination of basis images, are determined based on light source directions and the light map. By resampling the light map using the recovered light source directions, the object can be relit. In section 6, this resampling is elaborated.

2. Related Work

Relighting is not a new topic in computer graphics and considerable research has already been done in this field. The existing techniques can roughly be divided into two categories: geometry dependent relighting techniques and image-based relighting (IBR).

The geometry dependent techniques require a model of

[†] {vincentm,phil,frederik}@cs.kuleuven.ac.be

the object. This can be given or, using image-based modeling techniques, a geometric approximation of the object can be obtained. Usually, geometry dependent methods derive surface properties of the object, represented by a bidirectional reflectance distribution function (BRDF). With a geometric model and BRDFs it becomes possible to render the scene from new view points with novel lighting conditions by applying global illumination algorithms.

Loscos et al.⁹ use an approach based on the radiosity algorithm. From a set of photographs from a fixed view point, taken with different illumination settings, the diffuse albedo for each patch in the scene is estimated. Yu et al.¹⁷ use a large set of photographs taken from different view points with unchanged illumination to estimate the BRDFs of all surfaces. The Ward model¹⁴ is used as BRDF model so non-diffuse material properties can be extracted. Boivin et al.¹ use one picture to estimate the BRDFs. For all patches, the BRDF is approximated with an incremental and hierarchical technique. Malzbender et al.¹⁰ extract Polynomial Texture Maps or PTMs from a surface instead of BRDFs. Per texel, coefficients of a polynomial are stored. These are used to reconstruct the surface color under varying lighting conditions. Still, a coarse geometric model is needed to apply these PTMs for rendering.

A second category of relighting algorithms are the image-based relighting (IBR) techniques, which require no geometry at all. Wong et al.¹⁶ introduced the concept of an apparent BRDF. Based on a set of pictures, the BRDF of a pixel on the image plane is assembled. This technique was further improved to relight panoramas¹⁵. Lin et al.⁸ researched the dual of the Lumigraph⁵, in which the camera is fixed and a point light source is moved. Pictures are taken of an object while the light source is mechanically positioned at coplanar grid points. After resampling the data, the object can be relit. A geometrically independent bound on the grid resolution of the light source placement is derived as well. This bound is related to BRDF of the object. Nimeroff et al.¹¹ used a technique of combining images to relight a scene. Due to the linearity of illumination, objects can be relit by creating a weighted sum of basis images. The weights are calculated using steering functions. Debevec et al.³ describes a Light Stage in which an object can be placed. This device positions the light source at fixed positions using a gantry, while a camera takes pictures from a fixed view point. Using these basis images the object can be relit with any environment map. Evolutions of the Light Stage, Light Stage 2.0 and Light Stage 3.0, have been developed, the former speeding up the capturing process⁶, the latter enabling to relight objects with an arbitrary light map in real life.

3. Overview of our technique

Our proposed technique is an image-based relighting technique and is related to the Light Stage³. The Light Stage and all its evolutions use a gantry to move the light source to

known positions. This limits the sampling resolution of illuminant directions (i.e. the number of possible illuminant directions) and the size of the objects to be relit. We remove the limitation of the gantry, and put no restrictions on light source placement. Instead, the light source is moved freely around the object.

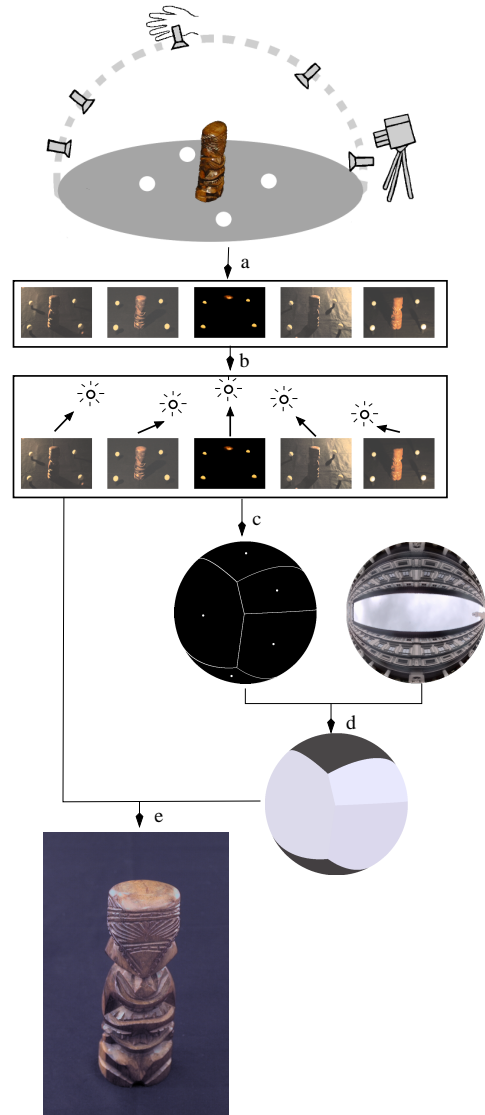


Figure 1: Schematic overview of our technique. With the correct setup the basis images are recorded (a). From the shading pattern of the diffuse white spheres the illuminant direction is estimated (b). The light source directions are plotted on the hemisphere and an angular Voronoi diagram is constructed. (c). For each Voronoi cell, the incident light from the light map is integrated (d) and this is used as scalar in the linear combination of the basis images to relight the object (e).

We work in three phases (figure 1). First we start with the acquisition of basis images. Basis images are pictures of the object taken from a fixed camera position. For each image, a hand-held light source is positioned freely, hence, the illumination per basis image is different. The illuminant direction is unknown but will be estimated since it is required for re-lighting. Next to the object, four diffuse white spheres are placed. These spheres are visible in the basis images as well.

In the second phase we use the parts of the basis images showing the diffuse spheres, to estimate the illuminant direction. Since diffuse spheres and a single light source are used, we can derive the illuminant direction from the shading pattern of the spheres.

Finally, we relight the object with an incident light field. This light field is usually given as an environment map. For each basis image, a weight needs to be calculated. This weight will be determined by the illuminant direction of that basis image and the incident light to be used to relight the object. We plot all light source directions on a hemisphere. An angular Voronoi diagram is then constructed using the plotted directions as sites. The weight for a basis image is the integration of all incident light in the Voronoi cell related to that basis image.

In this paper, wherever we refer to the reflectance field, we mean the non-local reflectance field as defined by Debevec et al³.

4. Data Acquisition

Figure 2 shows the complete setup for data acquisition. A digital camera, a hand-held light source, an object to be relit and 4 diffuse white spheres are needed.




Figure 2: A small statuette is illuminated by a hand-held light source and continuously photographed, resulting in several pictures with different illumination.

The white diffuse spheres are placed around the object and are visible by the camera as well. The spheres will be used to estimate the light source direction, since the hand-held light source will be moved freely but the light direction still needs to be known. The whole system is placed in a

darkened room, in which the hand-held light source is the only source of illumination.

The setup can be done in as few as a couple of minutes. Typically, 400 to 500 pictures were taken. The photo shoot is easy and is realized in 25 to 30 minutes, we use a Canon EOS D30 camera. The camera is controlled by a computer and takes a photograph every 4 seconds. The hand-held light source is moved freely but held still every time a picture is taken. The light source movement can be irregular, so large sections of the hemisphere could be under-sampled. To prevent this from happening, it might be a good idea to end the photo shoot with a series of pictures in a regular pattern emulating the gantry from the original Light Stage.




Figure 3: The Free-form Light Stage in action. The hand-held light source is moved freely while pictures are taken.

Care has to be taken that the spheres do not affect the illumination of the object, since they are not part of the object to be relit and their influence on the illumination of the object should not be noticeable. Different problems could arise:

- *The spheres are in front of the object.* In the resulting image of the relit object, the spheres will block the view of the object as well. However, the spheres will be correctly relit, as they are part of the scene.
- *The spheres are reflected in a highly specular object.* In the image of the relit object the reflected and relit spheres will be noticeable.
- *A sphere casts a shadow on the object.* The object can be relit with a light map, but an unwanted blurred shadow of the sphere will be seen in the image of the relit object.
- *The object casts a shadow on one of the spheres.* This will not have an effect on the image of the relit object. It will however complicate the extraction of the illuminant direction.

These effects can be avoided by placing the spheres with caution. By placing the spheres at the side and behind the object, shadows of the spheres will be cast on the back of

the object, thus unseen and reflections will be hardly noticeable. Basis images with shadowed spheres can be rejected for further use, since no decent illuminant direction can be estimated.

The spheres are best placed as far as possible from the object, reducing the interaction and casting of shadows to a minimum. However, the farther the spheres are set from the object, the more the camera needs to be zoomed out to capture all the spheres as well as the object. This would result in only a small set of pixels in each basis image representing the object. Although modern digital cameras have a very high resolution, still a lot of detail would not be captured. This can be overcome by using two synchronized cameras. The first camera is zoomed out to capture the whole set of spheres and these images are used to estimate the light source direction. The second camera is focused on the object alone and these pictures can be linearly combined.

The light source used, is a hand-held calibrated Xenon light source, emitting diffuse light and is small in area ($10\text{cm} \times 10\text{cm}$). This type of light source is adequate for our assumptions (i.e. a non-directional point light source). In Figure 3 we see the Free-form Light Stage in action.

The linear combination of photographs should be done using radiance values. The camera returns pixel values, the result of a non-linear transformation applied to the incoming radiance. We take low dynamic range (LDR) pictures but minimize the saturation by tuning the aperture and exposure time of our camera. We extract the response curve of our digital camera with the technique proposed by Debevec et al.⁴. By transforming our LDR pictures into high dynamic range (HDR) pictures using the inverse response curve, we obtain the radiance values.

5. Light Source Direction Estimation

Once the pictures are taken, we will analyze each picture to find the illuminant direction. Light source estimation from real images is a well researched problem in computer vision. Research has been done extensively on problems of determining 3D shape from two-dimensional shading. An extensive overview can be found in Zhang et al.¹⁸.

We started from the approach outlined by Pentland¹², and changed the idea to suit our needs. They were able to determine shape from shading even in cases where illumination direction, illuminant strength, and surface reflectivity are not known. By making assumptions about changes in surface curvature and the position of surface points, least square methods solve their model for illuminant direction and surface orientation. A follow-up paper by Lee and Rosenfeld⁷ essentially uses the same assumptions, but is based on statistical methods to determine the light direction and then estimates shape orientation in the light source coordinate system.

In our algorithm, we do not need shape recovery, but need




Figure 4: Conceptual representation of the photographed sphere. The sphere is located somewhere in the scene. Its perspective projection is displayed on the image plane. By computing the normal direction \bar{N}_p for each 'sphere pixel' p , we are able to obtain the illuminant direction.

an accurate method to determine the illuminant direction. Note that we do not try to recover the light source position. Only the direction as measured from the center of the scene will be recovered. Wherever most methods are not accurate enough because they are based on assumptions about unknown geometry and unknown illuminant direction, we wanted an accurate method standing on known geometry.

Based on the model of Lambert shading, extraction of an illuminant direction is possible if the exact normals are known of a sphere. Those normals can be computed easily if the center of the sphere is known. The algorithm outlined here searches the illuminant direction, given a calibrated camera and the exact sphere radius. Optimization is used to find an illuminant direction after computing the sphere center and the normals for its surface points (figure 4).

A numerical algorithm is necessary for extraction of the sphere position with respect to the camera. If the center is wrongly estimated, all normals will be false and our precision on illuminant direction decreases. As can be easily verified in figure 4,

$$\|C\|^2 - R^2 - (\bar{s}_i \cdot \bar{C})^2 = 0, \quad (1)$$

where \bar{s}_i are the silhouette directions. \bar{C} is the sphere center and R its radius.

To find the silhouette directions \bar{s}_i , we singled out the region of the image where the sphere is displayed, by using a matte. Afterwards, a Canny Edge detection algorithm searches all the visible edge pixels in that region. Using the calibration parameters of the camera, the silhouette directions \bar{s}_i can be computed from these edge pixels. A good es-

timation for \bar{C} can then be found by non-linear optimization techniques, minimizing equation 1.

The computation of the surface normals \bar{N}_p for each point of the sphere is quite straightforward now. Again, we use the region with the sphere data and make a list of those pixels containing useful intensity information. However, due to inter-reflections in our scene, the spheres might contain some additional information. So it is useful to position the object and spheres on black velvet in a dark non-reflective environment. It prevents color bleeding from the scene on the spheres. By intersecting the sphere with \bar{o}_p , \bar{P} and \bar{N}_p can be computed. The normal \bar{N}_p is the desired normal direction for each pixel in the list.

Finally, the normals and shading information together can be used to find the illuminant direction. It is based on the principle of Lambert's cosine law which states that the reflected intensity of light I on a surface is proportional to the cosine of the incident angle of the light reaching the surface.

$$I_p = \rho I_L (N_p \cdot L), \quad (2)$$

where I_L is the point light source intensity, ρ the material's diffuse hemispherical reflectivity, N_p resp. L are the surface normals and the illuminant direction.

Least square optimization techniques can now solve equation 2 for the illuminant direction L using the intensity I_p and normal N_p data from all pixels containing shading information.

Before we applied our technique to real pictures, we tested this procedure on a virtual environment. When the setup consists of a point light source at infinity, which is mathematically correct according to equation 1, the estimated direction differs at most 1 degree from the exact direction. For the Free-form Light Stage, a small area light source is used at approximately 2 meters distance. The accuracy decreases with 2 to 3 degrees as can be seen in table 1. When our technique is applied to real scenes, experiments showed that the error still increases due to noise, but does not become larger than 7 degrees.

In practice, to obtain the best results for the Free-form Light Stage we opt for placing 4 spheres symmetrically around the center of our real scene (fig. 5). The illuminant direction is computed for each sphere separately. As is illustrated in figure 5, the best fitted intersection is used to find the illuminant direction with respect to the center of the scene.

6. Relighting Objects

After the data acquisition and illuminant direction estimation, the third and final stage is the actual relighting. To relight the object we use an incident light field, given as an environment map. We relight our objects with light coming from the upper hemisphere.

ϕ/θ	0	10	20	30	40	50	60	70
0	1.0	1.5	1.7	1.7	1.2	0.5	1.1	1.8
20	1.1	0.7	1.4	2.5	1.9	1.8	0.8	0.7
40	1.1	2.8	2.4	3.1	2.8	2.5	1.9	1.0
60	1.0	2.5	2.4	2.7	3.0	2.8	3.0	1.5
80	1.1	1.9	2.4	2.5	2.8	2.7	2.5	2.8
100	1.1	1.7	2.8	2.2	1.7	2.0	2.8	3.4
120	1.1	1.3	1.3	1.4	1.2	3.2	3.2	4.1
140	1.1	0.7	0.6	1.5	5.1	2.7	3.4	3.2
160	1.1	0.3	0.5	0.8	2.0	2.7	2.7	2.5
180	1.1	0.3	1.6	2.2	2.7	3.6	0.7	4.2

Table 1: Degree error table: The error is plotted for a virtual setup as shown in figure 5. θ is the downwards direction into the image-plane orientation. Φ is the azimuthal angle, 0 degrees being the camera direction. The estimated illuminant direction differs at most three to four degrees from the exact direction.

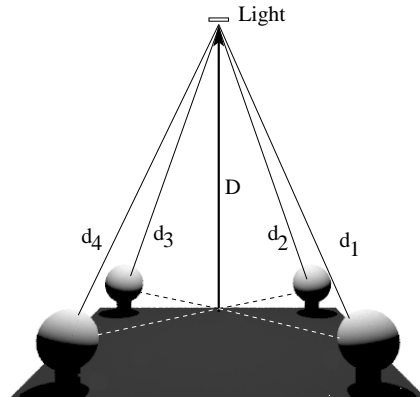


Figure 5: Estimating the Illuminant Direction with 4 spheres. \bar{D} is calculated using $\bar{d}_1, \bar{d}_2, \bar{d}_3$ and \bar{d}_4 .

To relight the object correctly, all the incident light of the environment has to be used. For each picture we have an associated illuminant direction. We will distribute all the incident light of the light map over these light source directions. This distribution will then determine the weights for each picture in the linear combination, resulting in the final picture.

In figure 6, a mirrored ball in an environment is displayed, accompanied with only the upper hemisphere of the environ-




Figure 6: Left: a mirror ball placed in the Uffuzi gallery. Right: the upper hemisphere of that environment. Light map from <http://www.debevec.org/Probes>




Figure 8: Left: The hemisphere of Figure 6 with the angular Voronoi cells overlaid. Right: The resampled incident light map.

ment. In this and following images the hemisphere is projected on a disk.

Every pixel in the light map represents a small solid angle with some incident light coming from that direction and we want to associate each pixel with the closest sampled illuminant direction. This can be done in the following way: every extracted light source direction is plotted on the hemisphere. From the resulting set of points, an angular Voronoi diagram is constructed, using the angle between directions as the distance criterion.

This angular Voronoi diagram construction is displayed in figure 7 (a) and (b) for three illuminant directions. An angular Voronoi diagram with 50 light sources is displayed in Figure 7 (c).

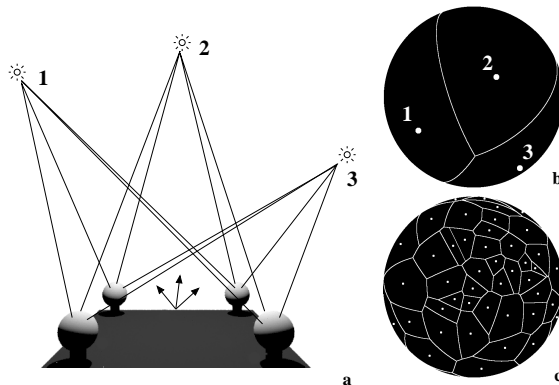


Figure 7: Construction of the Voronoi diagram in on the hemisphere, based on the estimated illuminant directions.

This gives for every illuminant direction a Voronoi cell, representing all incident light directions closest to that illuminant direction. We now integrate all the incident light over each Voronoi cell (figure 8). The resulting radiance value will be used as the weight for the accompanying basis image, illuminated from the corresponding light source direction.

Constructing an angular Voronoi diagram is independent of the parameterization used to represent the incident light map. Any hemispherical parameterization can be used. We prefer the concentric map parameterization, proposed by Shirley et al.¹³, since it preserves adjacency. The preservation of adjacency simplifies the integration of incident light on the Voronoi cells. In our implementation, the weight for a basis image is the sum of all intensities in the corresponding angular Voronoi cell. Since every pixel represents a different solid angle, care has to be taken that each pixel is weighted appropriately.

The Voronoi diagram brings an elegant solution to the non-uniform distribution of light source directions over the hemisphere, which results from the uncontrolled and free movement of the hand-held light source.

A comparison with the Light Stage is interesting: if the sampled illuminant directions are the same as when using a gantry (i.e. a regular sampling), then the same weights will be calculated as with the original Light Stage.

In our examples, only the upper hemisphere of incoming light is used, since we place the object on the floor, and illumination only comes from above. Extension to a full sphere of light directions is trivial (e.g. place the object and spheres on a tripod).

7. Results

We have experimented with a range of different objects. Results can be seen in figure 11. A Hawaiian statuette and a set of toy soldiers are relit with four different light maps. For the Hawaiian statuette, 400 basis images were taken, for the toy soldiers 500 images were used. In figure 11 (a) an artificial light map is applied. Notice the colored shadows (e.g. around and under the chairs of the soldiers). Red light clearly comes from the back, while yellow light is coming from the front. This can be seen clearly on the statuette and on the shoulders of the soldiers. In (b), the objects are relit using the upper hemisphere of the environment map in a

room with four ceiling lights. Soft shadows can be seen as well as highlights. In (c), the bright windows on the right of the light map can be noticed as a reflection (e.g. on the right side of the statuette) and leave elongated soft shadows (e.g. on the set of soldiers). In (d), a more homogeneous light map was used, resulting in very soft shadows and very few highlights.

We also tested the influence of the number of basis images to the quality of relighting. The object was relit twice with the same light map, but with a different set of basis images. From a set of 500 basis images, a subset of 50 randomly selected basis images was used to produce figure 9 (a). The complete set was used to generate figure 9 (b).




Figure 9: Comparison between relighting with a set of 50 basis images (a) and a set of 500 basis images (b). Both sets are random sampled illuminant directions

In figure 9, the artificial light map from figure 11 (a) was applied. When a reduced set of images is used, the colored soft shadows are not reproduced faithfully. This is because the small set holds insufficient information to relight the scene realistically. Combining only a few basis images in a weighted sum, allows only coarse interpolation. Therefore, soft shadows and sharp highlights can only be accomplished using an extensive set of images.

As an application, our technique can be used to place a real object in a virtual scene. During photographing, position the object on a surface that has approximately the same BRDF as the surface on which the object will be placed on, in the virtual scene. This will ensure that the correct amount of light will be reflected from the surface to the object. The reflectance field of the object is then captured using our presented technique. Remove the object in the

setup and capture another reflectance field of just the surface. Using RenderPark, a global illumination software package (www.renderpark.be), we create an environment map at the approximate position where the object will be placed in the virtual scene. We can then relight the two reflectance fields with this rendered environment map and, using a differential rendering technique², the real object is consistently illuminated and merged with a rendered image of the virtual scene. In figure 10, the top image is the original rendering. The image in the middle is the augmented virtual scene and the bottom image zooms in on the photographer and his camera. The shadows of the soldiers give strong visual cues to link the real and virtual objects together. A similar technique can also add real objects to a photograph of a real scene.



Figure 10: On the top a rough rendering of a table with chart and globe is displayed. At the bottom the set of soldiers is added. Notice the correct shadows in the zoom-in.

We applied our technique on other objects, including a human face, and experimented with other incident light maps. These results, together with some animations and the used environment maps, can be found on <http://www.cs.kuleuven.ac.be/~graphics/CGRG.PUBLICATIONS/FFLS>.

8. Conclusion and Future Work

We have presented a technique to capture the reflectance field of a real object using a hand-held free moving light source.

Future research includes better sampling of the incident light map. In our current implementation, we used a box filter over each Voronoi cell, but Gaussian filters might produce better results.

We have only used one camera position. Moving the camera while moving the light source as well might enable us to relight a real object and render it from any view point. This total freedom might be interesting to pursue in future research.

Since we do not use a gantry, or any other device to place light sources at known positions, the size of the object is only limited by the ability to move the light around the object. The relighting of large objects (e.g. a car) therefore becomes a possibility. However, positioning the light source at all possible positions could be more challenging. Except for the light movement, bigger diffuse spheres for light source detection would be needed, since our small spheres would be hardly noticeable in a photograph with large objects. Other illuminant direction recovery techniques could also be researched.

9. Acknowledgments

Many thanks to Pieter Peers and Frank Suykens for the countless useful suggestions. We wish to thank Paul Debevec for the environment map of the Uffizi gallery (<http://www.debevec.org/Probes>). This research was supported by FWO Grant #G.0354.00.

References

1. Samuel Boivin and Andr e Galalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, Annual Conference Series, pages 107–116. ACM SIGGRAPH, Addison Wesley, August 2001. 2
2. Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In Michael Cohen, editor, *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pages 189–198. ACM SIGGRAPH, Addison Wesley, July 1998. 7
3. Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In Kurt Akeley, editor, *SIGGRAPH 2000, Computer Graphics Proceedings*, Annual Conference Series, pages 145–156. ACM SIGGRAPH, Addison Wesley, July 2000. 2, 3
4. Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In Turner Whitted, editor, *SIGGRAPH 97 Conference Graphics Proceedings*, Annual Conference Series, pages 369–378. ACM SIGGRAPH, Addison Wesley, August 1997. 4
5. Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. In Holly Rushmeier, editor, *SIGGRAPH 96, Computer Graphics Proceedings*, Annual Conference Series, pages 43–54. ACM SIGGRAPH, Addison Wesley, August 1996. 2
6. Tim Hawkins, Jonathan Cohen, and Paul Debevec. A photometric approach to digitizing cultural artifacts. In *In 2nd International Symposium on Virtual Reality, Archaeology, and Cultural Heritage, Glyfada, Greece, November 2001.*, 2001. 2
7. C. H. Lee and A. Rosenfeld. Improved methods of estimating shape from shading using the light source coordinate system. *Artificial Intelligence*, 26:125–143, 1985. 4
8. Zhoulin Lin, Tien-Tsin Wong, and Heung-Yeung Shum. Relighting with the reflected irradiance field: Representation, sampling and reconstruction. *IEEE Computer Vision and Pattern Recognition*, 1:561–567, December 2001. 2
9. C eline Loscos, George Drettakis, and Luc Robert. Interactive virtual relighting of real scenes. In *IEEE Transactions on Visualization and Computer Graphics*, volume 6(4), pages 289–305. IEEE Computer Society, 2000. 2
10. Tom Malzbender, Dan Gelb, and Hans Wolters. Polynomial texture maps. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, Annual Conference Series, pages 519–528. ACM SIGGRAPH, Addison Wesley, August 2001. 2
11. J Nimeroff, E Simoncelli, and J Dorsey. Efficient re-rendering of naturally illuminated environments. In *Eurographics Rendering Workshop 1994*, Darmstadt, Germany, June 1994. Springer-Verlag. 2
12. Alex P. Pentland. Finding the illuminant direction. *Optical Society of America*, 72(4):448–455, April 1982. 4
13. Peter Shirley and Kenneth Chiu. A low distortion map

- between disk and square. *Journal of Graphics Tools: JGT*, 2(3):45–52, 1997. 6
14. Gregory J. Ward. Measuring and modeling anisotropic reflection. In Edwin E. Catmull, editor, *SIGGRAPH 92, Computer Graphics Proceedings, Annual Conference Series*, pages 265–272. ACM SIGGRAPH, July 1992. 2
 15. Tien-Tsin Wong, Pheng-Ann Heng, and Chi-Wing Fu. Interactive relighting of panoramas. *IEEE Computer Graphics and Applications*, 21(2):32–41, March/April 2001. 2
 16. Tien-Tsin Wong, Pheng-Ann Heng, Siu-Hang Or, and Wai-Yin Ng. Image-based rendering with controllable illumination. In Julie Dorsey and Philipp Slusallek, editors, *Eurographics Rendering Workshop 1997*, pages 13–22. Eurographics, Springer Wien, June 1997. 2
 17. Yizhou Yu, Paul Debevec, Jitendra Malik, and Tim Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs from. In Alyn Rockwood, editor, *SIGGRAPH 99, Computer Graphics Proceedings, Annual Conference Series*, pages 215–224. ACM SIGGRAPH, Addison Wesley, August 1999. 2
 18. Ruo Zhang, Ping sing Tsai, James Edwin Cryer, and Mubarak Shah. Analysis of shape from shading techniques. *IEEE Computer Vision and Pattern Recognition*, pages 377–384, June 1994. 4

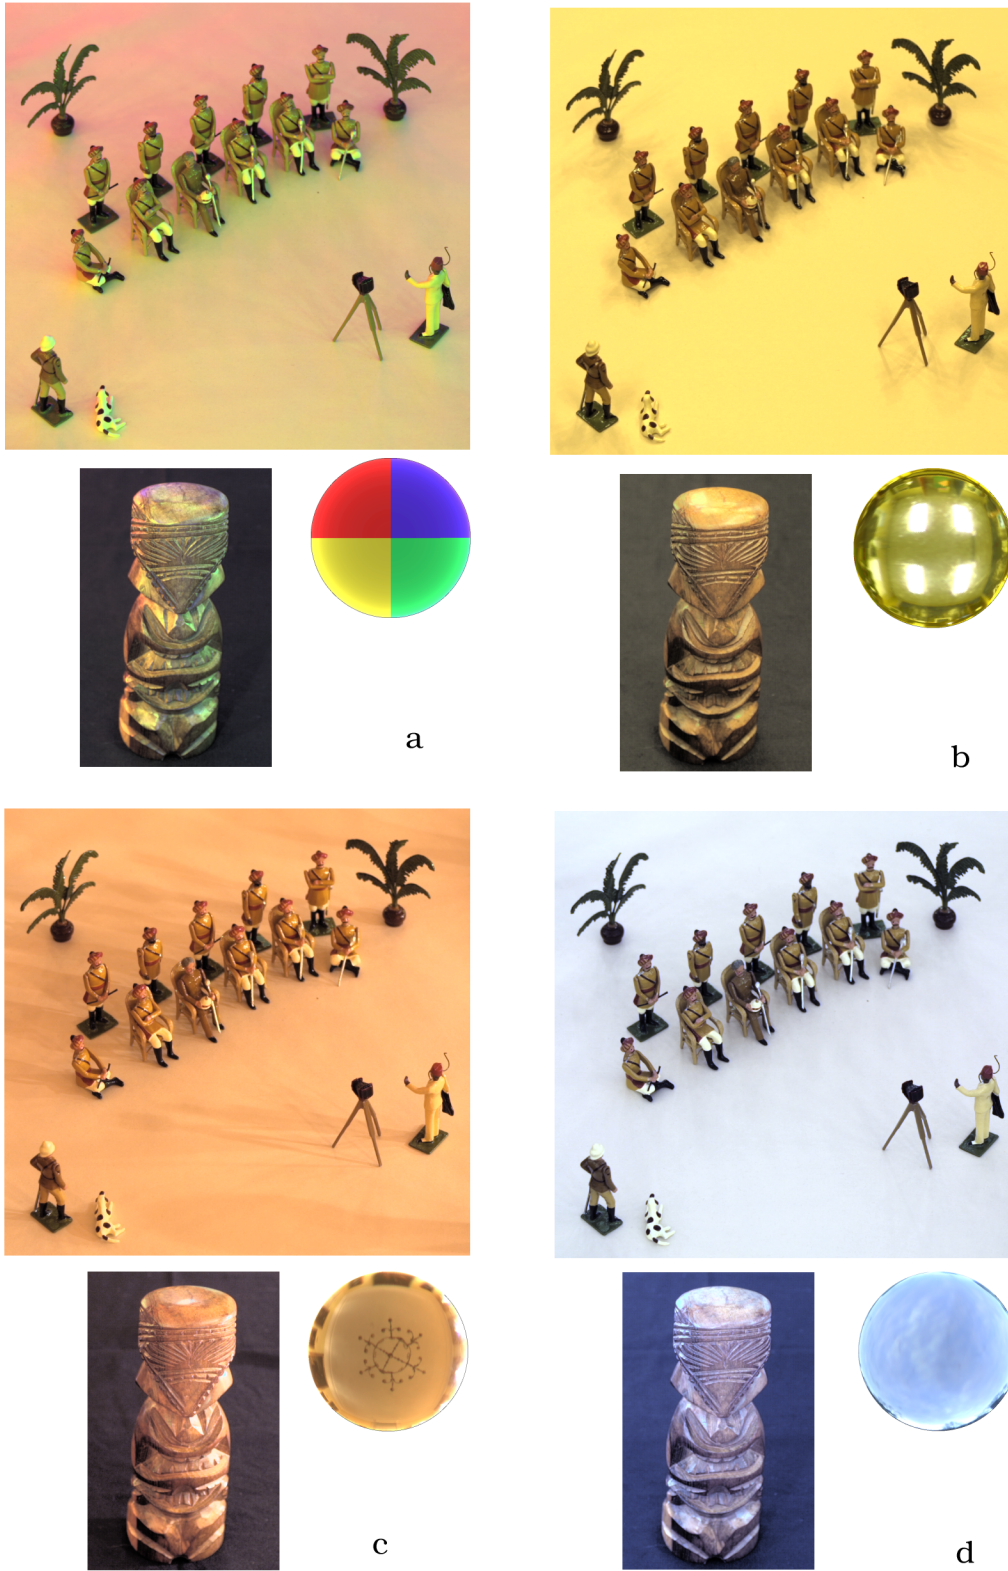


Figure 11: A Hawaiian statuette and a set of toy soldiers are relit with our technique. Each couple of images is accompanied with the applied incident light map.