

Capturing and Rendering With Incident Light Fields

J. Unger,¹ A. Wenger,² T. Hawkins,² A. Gardner² and P. Debevec²

¹ Linköping University Norrköping Visualization and Interaction Studio, Sweden

² University of Southern California Institute for Creative Technologies, United States

Abstract

This paper presents a process for capturing spatially and directionally varying illumination from a real-world scene and using this lighting to illuminate computer-generated objects. We use two devices for capturing such illumination. In the first we photograph an array of mirrored spheres in high dynamic range to capture the spatially varying illumination. In the second, we obtain higher resolution data by capturing images with an high dynamic range omnidirectional camera as it traverses across a plane. For both methods we apply the light field technique to extrapolate the incident illumination to a volume. We render computer-generated objects as illuminated by this captured illumination using a custom shader within an existing global illumination rendering system. To demonstrate our technique we capture several spatially-varying lighting environments with spotlights, shadows, and dappled lighting and use them to illuminate synthetic scenes. We also show comparisons to real objects under the same illumination.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Light Fields, Illumination, Image-Based Rendering, Reflectance and Shading, Image-Based Lighting

1. Introduction

Lighting plays a key role in the realism and visual interest of computer generated scenes, and has motivated many techniques for designing and simulating complex forms of illumination for computer-generated objects. Recently, techniques have been proposed that capture light from the real world to be used as the illumination for computer-generated scenes. Using images of real-world illumination has proven useful for providing realistic lighting as well as for effectively integrating computer-generated objects into real-world scenes. So far, most techniques for capturing real-world illumination acquire a lighting environment from a single point within a scene. While such a measurement records the directionally varying illumination from different light sources and surfaces in a scene, it does not capture the spatially varying illumination in the scene, i.e. how the light varies from one location to another. In lighting design and cinematography, spatially varying illumination such as cast shadows and shafts of light plays an important role in the visual makeup of artistically designed scenes. This creates a need to be able to capture the spatially as well as directionally varying illumination within a scene.

Recently, the light field and Lumigraph concepts pre-




Figure 1: *The high fidelity incident light field capturing device. The Uniq UP-1030 camera is mounted onto the translation stage, allowing it to move in the plane and sample the incident light field.*

sented techniques for recording the spatially varying appearance of objects and environments. The techniques work by recording a two-dimensional array of images across a surface, and can extrapolate the scene's appearance to a three-

dimensional volume by tracing rays to the original capture surface.

In this work, we leverage these techniques to the capture of spatially varying illumination, and thus produce a technique which allows us to illuminate computer-generated objects with captured light fields and Lumigraphs. We call this capture of spatially varying illumination an incident light field.

In a manner similar to the capturing of traditional light fields, we constructed two devices to acquire an incident light field by taking angular lighting measurements over a two-dimensional plane. We then use this captured light volume to illuminate a synthetic scene.

2. Background and Related Work

The work we present in this paper is based on two principal ideas: first, that light in space can be captured as a two-dimensional array of images within an environment, and second, that light captured in a real-world environment can be used to illuminate computer-generated objects.

The way in which illumination varies within space was described by Adelson and Bergen¹ as the *Plenoptic Function*. They proposed this seven-dimensional function in the form $P = P(\theta, \phi, \lambda, t, V_x, V_y, V_z)$, where P is defined as the radiance arriving at a point (V_x, V_y, V_z) in the direction (θ, ϕ) at time t with wavelength λ . They noted that this function contains every omnidirectional image of the world which can be recorded at any time. Fixing time and discretizing wavelength to three spectral integrals for red, green, and blue, the five-dimensional function of $(\theta, \phi, V_x, V_y, V_z)$ contains every omnidirectional image of every point in a space. Levoy and Hanrahan¹⁵ and Gortler et al.⁹ noted that since the radiance along a ray is constant in unoccluded space, a 2-dimensional array of images on a plane called a *Light Field* or *Lumigraph* can be used to reconstruct a 5D plenoptic function from a 4D dataset. Ashdown² presented *Near Field Photometry*, a method for capturing exitant light fields using CCD cameras. In our work we use the Light Field technique to capture the incident illumination within a space, and to extend this information to the remainder of the space. To make maximum use of our spatial sampling, we use the depth correction technique described by Gortler et al.⁹ to project and focus our captured light onto an approximate geometric model of the scene. Heidrich et al¹¹ mentions the idea of using Light Fields to illuminate computer generated scenes. Our work builds on the same idea, but captures real world lighting and utilizes global illumination techniques to illuminate synthetic scenes. Heidrich et al¹⁹ used pre-computed light field information, *Canned Light Sources*, for illuminating virtual objects.

Using images of incident illumination to produce realistic shading was introduced as the process of environment mapping by Blinn and Newell³. Miller and Hoffman¹⁷ captured

environment maps from the real world by photographing a reflective sphere placed in the environment, and re-mapped the sphere's image onto the surface of an object to show how the object might reflect the light from that environment. They showed that blurring an environment map before applying it to the surface of an object could be used to simulate a variety of object reflectance properties.

Debevec et al⁴ used omnidirectional measurements of incident illumination as light sources within a global illumination rendering context, and showed that high dynamic range photography techniques (in which a series of differently exposed digital images is combined into a single radiance image covering the entire dynamic range of the scene) is useful for capturing the full range of light encountered in the real world. Used directly, these illumination capture techniques record light only at a single point in space, and thus do not capture how light directions, colors, and intensities change within an environment.

Sato et al.¹⁸ and Debevec et al.⁴ took initial steps toward recording spatially-varying illumination by projecting the captured illumination onto a 3D model of the environment obtained through stereo¹⁸ or photogrammetry⁴. However, the technique would not capture any significant changes in incident illumination such as harsh shadows or spotlights. For a synthetic scene, Greger et al.¹⁰ computed its *Irradiance Volume*, a 3D set of diffuse environment maps located in an evenly spaced lattice within the scene. With these pre-computed irradiance measurements, they interactively rendered a new diffuse object moving in the scene with nearly correct incident illumination. Looking at time instead of space, Koudelka et al.¹⁴ recorded time-varying incident illumination, used to capture and reproduce the illumination from a moving light source. Yang et al.¹³ and Goldlücke et al.⁸ used digital video cameras to capture time varying light fields.

This paper begins with this body of previous work to record spatially varying real-world illumination which includes harsh shadows and directional light sources, and uses such a captured field of illumination to realistically render new objects into this spatially-varying lighting.

3. Incident Light Field Parametrization and Sampling

A light field¹⁵ describes the distribution of light in a static scene with fixed illumination. Our interest in the light field lies in its application as lighting information for rendering virtual objects with real world light. As in Debevec et al.⁵ we refer to such a lighting dataset as an *incident light field*, or ILF. We note that an ILF is fundamentally the same as a standard light field; the term is meant only to emphasize that an ILF is meant to capture the full dynamic range and directionality of the incident illumination, which is rarely done for light fields meant for direct viewing. Several light field parameterizations have been proposed, including the two-plane




Figure 2: Incident Light Field Parametrization. The incident light field is defined by a reduced plenoptic function $P = P(\theta, \phi, u, v)$, where (u, v) define the location of \vec{r}_o in a plane and (θ, ϕ) define the direction \vec{r}_d towards the incident light with respect to the plane normal \vec{n} .

parameterizations presented by Gortler et al. ⁹ and Levoy et al. ¹⁵. Since we are specifically interested in capturing incident illumination, we base our light field parametrization on a point (u, v) on a particular capture plane Π the set of incident illumination directions (θ, ϕ) incident upon Π . For the unoccluded volume above the capture plane, the incident light field $P(\theta, \phi, u, v)$ can be used to extrapolate the incident illumination conditions within the volume.

In this work we only consider the incident light coming from the hemisphere above the normal \vec{n} of the plane Π . Our light field function is therefore over the domain of:

$$P = P(\theta, \phi, u, v), \text{ where } \begin{cases} -\pi \leq \theta \leq \pi \\ 0 \leq \phi \leq \frac{\pi}{2} \end{cases}$$

ϕ and θ are the direction with respect to \vec{n}
 u and v are the coordinates in the plane Π

The above parametrization describes the incident light field as a collection of rays $\vec{R}(s) = \vec{r}_o + s \cdot \vec{r}_d$, where \vec{r}_o lies in the plane Π with coordinates (u, v) and where \vec{r}_d is defined by (θ, ϕ) as in Figure 2.

The continuous incident light field is discretized by sampling the hemisphere of incident light with a fixed angular resolution at locations (i, j) on an $n \times m$ regular grid in the plane Π . The above parametrization allows us the capture a discrete light field as a 2D array of light probes, where a single (θ, ϕ) image captures the directional information and the position of the light probe in the 2D array specifies the (u, v) coordinates in the plane.




Figure 3: The mirror sphere array. The incident light field is sampled by assembling a high-dynamic range image of the light reflected from the plane by an array of mirrored spheres.

4. Data Acquisition and Processing

For capturing incident illumination arriving at a single point in space, Debevec et al. ⁴ used a high dynamic range photography technique to capture the full range of luminance values in a real-world scene.

In this work we designed and built two incident light field capturing devices, each based on combining techniques and results from light field rendering and image-based lighting research. Our first apparatus is closely related to the high-dynamic range light probe acquisition technique introduced by Debevec et al. ⁴, whereas the second apparatus is based more directly on the methods employed in the light field rendering world. In the following sections we describe the experimental setup and the data processing required for each of the two ILF capturing devices.

4.1. Mirror Sphere Array

The first device we built extends the high-dynamic range light probe idea by placing a series of mirrored spheres on a regular grid and taking a series of differently exposed photographs to construct a high-dynamic range representation of the incident light field.

4.1.1. Experimental Setup

The capturing setup consists of an array of 12×12 1" diameter mirror spheres and a standard digital camera. As seen in Figure 3 the mirror spheres are mounted on a board which corresponds to the plane Π .

We found that the best incident light field results were obtained by capturing a nearly orthographic view of the array from a nearly perpendicular direction from the sphere plane.




Figure 4: Mirror spheres array data set. The processed incident light field captured with the mirrored sphere array. Each rectangle is a full latitude/longitude map corresponding to the incident light onto a single mirrored sphere.

We approximate the orthographic view by placing the Canon EOS D30 camera far away from the mirror sphere array and by using a telephoto lens.

A benefit of this technique is that the capture process is very quick, requiring a single high-dynamic range image. Unfortunately, a capture of this type yields either poor directional resolution or poor spatial resolution as the resolution of the camera is limited and must encompass the entire grid of mirror spheres.

4.1.2. Data Processing

The data processing consists of two steps. First the image series with different exposure times is assembled into a high-dynamic range image using techniques presented in Debevec et al. ⁷. Second, the individual light probe subimages are remapped into a latitude/longitude format. This is done with a custom software where the user outlines the four corner spheres of the mirrored ball array. The program then registers the light probe array image into a dataset similar to the one as shown in Figure 4.

Our parametrization of the incident light field only considers the hemisphere above the plane normal. A problem with this technique is that the spheres interreflect, even for some upward-facing directions near the horizon. Fortunately, the geometry of spheres placed next to each other produces an unoccluded field of view of 159.2 degrees, which is nearly the full hemisphere. To avoid artifacts in renderings, the lowest 10.4 degrees near the horizon should be discounted from the dataset. Figure 10 shows the artifacts that can result from the interreflections if they are not discounted

4.2. High Fidelity Capturing Device

The second device we built addresses the resolution issues of the mirror spheres array device. The device works by moving the camera around in the plane (as previously done in Isaksen et al. ¹² and Levoy et al. ¹⁶) and taking a single omnidirectional high dynamic range image for each position.




Figure 5: An ILF high-dynamic range image series, is a sequence of 16 images with increasing exposure time from left to right. The image series is used to reconstruct the high-dynamic range measurement of incident illumination for a specific location in the capturing plane.

The directional resolution is therefore determined by the resolution of the camera whereas the spatial resolution is determined by how many locations are sampled in the plane.

4.2.1. Experimental Setup

The high fidelity capturing device (Figure 1) consists of three main parts. The first is a computer controllable camera with a fisheye lens, the second is a translation stage onto which the camera is mounted, and the third is the software system that controls the device and processes the captured data.

The camera used is a Uniq UP-1030 2/3" single-chip digital color camera with a maximum resolution of 1024×1024 pixels. To capture the entire incident illumination hemisphere, we mounted a 185° Coastal Optics fisheye c-mount lens to capture the ILF images.

The translation stage which moves the camera in the plane is a MD-2 Dual Stepper Motor System from Arrick Robotics. The computer-controllable stepper motors make it possible to move the camera with high accuracy from one position to the next during the sampling process. At its widest settings, the translation stage can move the camera around in a $30'' \times 30''$ area with a spatial resolution of $0.01''$.

The translation stage moves the camera along a regular grid with a user defined resolution. At each position on the regular grid the camera takes a series of 16 images with increasing exposure time from $1/16384^{th}$ of a second up to 2 seconds (see Fig. 5), from which a high-dynamic range image can be reconstructed.

While the high fidelity device can capture an ILF at arbitrarily high spatial resolution, the drawback of this approach is that the capture time increases accordingly. With our software, a high dynamic range incident light field capture and data processing with spatial resolution of 30×30 light probes takes approximately 12 hours. This can be challenging as the light field to be captured must be constant the entire time, which generally restricts the technique to indoor environments.




Figure 6: Brightness falloff curves for the camera/lens system. The brightness falloff curve for each color channel red, green and blue. Incident light with a given intensity that falls on the fisheye lens along the optical axis produces a pixel value almost twice as bright as incident light with the same intensity that falls on the fisheye lens at a 90° angle from the optical axis. We calibrate for this effect in the data acquisition.

4.2.2. Data Processing

The data processing involves building the high-dynamic range image for each camera position, remapping the fisheye image to a latitude/longitude image in (θ, ϕ) , and a brightness falloff correction process. The latter process is necessary since the fisheye lens we use exhibits a noticeable brightness falloff for the pixels far from the optical axis. We calibrated the lens by moving a constant-intensity diffuse light source beginning at the optical axis and continuing to the periphery of the fisheye lens, taking images at twenty intervals along the way. Assuming a radially symmetric falloff pattern, we extrapolated these readings to produce the brightness falloff curve shown in Figure 6 which we used to correct our ILF images. Figure 7 shows the high-dynamic range, remapped and brightness falloff corrected incident light field data.

5. Rendering with an ILF

In this section we present our method for rendering synthetic objects illuminated by real world lighting using the captured incident light fields. The algorithm relies on global illumination for the light transport in the scene and was implemented as a custom material type in the Radiance²⁰ lighting simulation package.

Our parametrization of the incident light field maps any ray in world-space to an incident ray in the sampled incident light field. During rendering, the radiance contribution of a ray to the scene is calculated by spatially interpolating be-




Figure 7: High fidelity incident light field data. A 30×30 spatial resolution and 400×400 directional resolution incident light field captured by the high fidelity device.

tween adjacent light probes and directionally interpolating within each of the light probes using bilinear interpolation in a manner similar to that in Levoy and Hanrahan¹⁵.

5.1. The General Algorithm

The custom material type specifies the location and orientation of the incident light field plane in the scene and its spatial dimensions. Since our parametrization of the incident light field is restricted to the hemisphere above the normal \vec{n} of the light field capture plane Π , we place the synthetic scene in the volume above Π as in Figure 8.

In order for the shader to have an effect on the rendering we need to cause a ray-surface intersection. We therefore define additional geometry, in the simplest case a sphere encompassing the entire local scene, to which we apply the incident light field material type. We refer to the additional geometry as auxiliary geometry.

When a ray, $\vec{R}_i(s) = \vec{r}_{o_i} + s \cdot \vec{r}_{d_i}$, intersects with the auxiliary geometry it potentially corresponds to a ray that maps to an incident illumination value. Since all the light in the scene is assumed to be incident from the hemisphere above the plane, rays directed upwards with respect to the plane will map to an incident illumination value in the incident light field. Once a ray \vec{R}_i intersects with the auxiliary geometry it is traced backwards to the incident light field plane Π (see rays \vec{R}'_0 and \vec{R}'_1 in Figure 8). If the backwards-traced ray \vec{R}'_i intersects with the incident light field plane (see \vec{R}'_0 in Figure 8), the intersection point is used to perform the




Figure 8: Outline of the rendering algorithm . The directional sampling of the incident light field is determined by the direction of the ray \vec{R}_i , whereas the spatial sampling of the incident light field is given by the intersection of backwards traced ray \vec{R}_i' with the incident light field plane.

bilinear interpolation between adjacent light probes. If the back-traced ray does not intersect the plane (see \vec{R}_1' in Figure 8), its radiance contribution is approximated using the closest spatial sample in the incident light field data set. That is, we assume that the light field outside of the sampled area is spatially uniform and continues the samples at the boundary of the incident light field. The original ray directions, e.g. \vec{r}_{d_0} and \vec{r}_{d_1} , are used to directionally interpolate within the previously determined adjacent light probes around the intersection of \vec{R}_i' with the incident light field plane. Rays from the lower hemisphere do not contribute to the radiance in the scene (see \vec{R}_2 in Figure 8).

As the adjacent light probes used for the spatial interpolation are not in the same location on the incident light field plane, the radiance originating from a specific point in the lighting environment corresponds for each light probe to a slightly different direction. The closer the point of origin of the radiance to the incident light field plane the more the corresponding directions in the light probes deviate. To improve the sampling of the incident light field we introduce a depth correction process similar to that presented by Gortler et al. ⁹ into our rendering algorithm. Once we have determined which four incident illumination images need to be sampled for the spatial interpolation, the depth correction recalculates the directions such that the direction vectors converge at the intersection of \vec{R}_i with the auxiliary geometry (Figure 9). The lights in the scene are assumed to be directionally smooth.

5.2. Depth Correction

With depth correction, the auxiliary geometry becomes more than a means of intercepting the rendering pipeline, but rather it serves as depth information for the lighting environment captured by the incident light field. The closer the




Figure 9: The depth correction process. The incident illumination images we need to sample for the spatial interpolation are, as before, determined by the intersection of the backwards-traced ray \vec{R}_0' with the incident light field plane. Instead of just using the direction of \vec{R}_0' as the lookup for the directional interpolation, depth correction recalculates the directions such that they converge at the intersection of \vec{R}_0' with the auxiliary geometry.

auxiliary geometry approximates the geometry of the lighting environment that the incident light field was captured in, the more the depth correction will improve the directional sampling, and thus results in higher quality renderings with fewer bilinear interpolation artifacts.

6. Results and Discussion

To demonstrate renderings produced with the incident light field process, we produced renderings of collections of synthetic objects under our captured illumination measurements, and compare these renderings to digital images of real objects in these lighting environments. All of our renderings were generated using the Radiance system with our custom shader.

Figure 10 shows two synthetic cubes and a reflective sphere rendered using the incident light field captured with the mirror sphere array seen in Figure 4. The ILF consists of two local spotlights aimed into the center of the sphere array. The rendering appears to reproduce this incident illumination correctly, indicating shadows and highlights from each of the sources, as well as a sharp shadow across the orange box at the edge of the illumination. The spots of light around the horizon of the synthetic mirrored sphere exhibits some of the sphere interreflection artifacts mentioned in Section 4.

Figure 11 shows a real versus synthetic comparison for a lighting environment captured with the high-fidelity ILF capture device. The lighting consists of two principal sources: a yellow spotlight aimed at the center of the table and a blue light partially occluded by a nearby rectangular card.

Figure 11(a) shows a set of real objects set into this lighting environment, including a 3D print of a computer model of a Greek statue. With the real objects removed, we captured the illumination at a spatial resolution of 30×30 and a directional resolution of 400×400 pixels for a total of 144,000,000 captured incident light rays. We virtually recreated the configuration of objects using the computer model of the statue and spheres on a virtual table within Radiance, using the calibrated ILF camera and a reflectance standard to estimate the surface reflectances of the objects. We used a measured model of the room as the auxiliary geometry for depth correction, and created a rendering of the virtual objects within the ILF using a single ambient bounce and 10,000 indirect illumination rays per pixel for a total rendering time of 21 hours. The rendering produced using this technique appears in Figure 11(b). The rendering produced a very similar image to the real photograph, including the blue light on the statue's head which was extrapolated from the ILF based on the blue light falling to the statue's left on the table. Due to a slight misalignment of the ILF with the synthetic scene, the blue light on the statue's head begins slightly higher on the statue than in the photograph. To emphasize the importance of capturing the spatially-varying illumination, we also rendered the synthetic scene as if it were illuminated entirely by the lighting measurement taken by the ILF capture device in the blue area of the table seen in Figure 11(c). While the scene is still realistically illuminated, it clearly does not match the real photograph.

For our final example, we tested the spatial resolution of our capture device by placed a tree branch between a small yellow light source and the ILF capture device (Figure 7). Since the branch was placed close to the ILF plane, it produced particularly sharp shadows. We created a second scene with the statue on a table under this illumination and photographed it as seen in Figure 12(a). Then, as before, we recreated the scene and camera viewpoint within the computer and illuminated the scene with the captured incident light field. The resulting rendering in Figure 12(b) is again largely consistent with the real-world photograph, but it clearly does not fully reproduce the high spatial variation of the dappled lighting from the tree branch. The reason is that the ILF's spatial resolution of 30×30 is far less than what is needed to fully capture the details of the lighting for this scene. Nonetheless, the ILF rendering of the scene is still far closer to the photograph than rendering the scene with a single measurement of scene illumination as in Figure 12(c).

7. Future Work

The capabilities and limitations of the current capturing and rendering system suggest several improvements for future versions of the system in the capture part as well as the rendering part.

The incident light field capture time for the high fidelity




Figure 10: Rendered scene with incident light field data from the mirror sphere array. This scene is rendered with the incident light field dataset from Figure 4. The inter-reflection artifacts from the spheres in the capturing device can be seen at the horizon of the mirror sphere in the image.

device is too high for the device to have any practical application, for example on a movie set. The capture time could be improved by using a better camera that has a good low light sensitivity and a high dynamic range, thus reducing the number of images required for the high-dynamic reconstruction. The good low light capabilities would reduce the exposure time needed to reconstruct the lower end of the high-dynamic range image. The capture time could be further improved by adaptively sampling different areas of the plane. For instance an irradiance map of the plane could be used to find areas with high spatial variation. The higher frequency areas could be sampled less than areas with higher variation. Thus reducing the overall number of samples need to represent the information of the incident light field.

The rendering part of the system could be improved by recovering depth information from the 2D array of light probes. This depth information can then be used instead of the rough scene approximation geometry to yield a much better depth correction. We could further use the depth information for a more efficient spatial interpolation between light probes. Using the depth information we can determine when light sources become occluded and can possibly reconstruct high frequency spatial variations in the incident light field. Working with depth information also relaxes the assumption that all the incident light originates from a distant scene. The only requirement remaining is the space described by the incident light field must be free of occluders. The smarter interpolation further has the potential of reducing the number of samples needed to reconstruct the continuous incident light field, which in turn would even further reduce the capture time.

Rendering quality and time could be drastically improved

by applying importance sampling to diffuse bounces based on the intensity values in the incident light field, instead of simply picking a random direction. This would produce less noisy images in less rendering time.

Looking at the data size of the captured incident light fields, data compression might be very useful when handling incident light fields. Capturing the single light probes as compressed images might reduce the rendering time because of the reduced amount of data transfer.

Another interesting avenue to pursue would be the application of incident light fields to lighting reproduction. Imagine a lighting reproduction apparatus that is capable of producing shadows moving across an actor's face or leaving one actor in shadow while another is illuminated by a street light. Incident light fields hold the needed information to reproduce such complex lighting environments. Recent work in the area of lighting reproduction Debevec et al. ^{5 6} made first steps in the direction of realistic compositing and relighting of live performances. The above approaches would benefit significantly from the additional information provided by incident light fields.

8. Conclusion

In this paper we presented a novel approach for rendering synthetic objects with spatially-varying real world illumination by fusing ideas from image-based lighting with ideas from light field rendering. We designed and built two devices to capture incident light fields, a simple low fidelity device using an array of mirrored spheres, and a high fidelity capture using a computer-controlled digital camera with a fish-eye lens mounted to a computer-controlled translation stage. We further showed how to render synthetic objects illuminated by the captured incident light fields using modifications to standard global illumination techniques. We believe our approach offers potential for creating visually interesting and highly realistic images as well as for integrating computer-generated objects into complex real-world environments.

9. Acknowledgements


We gratefully acknowledge Brian Emerson and Mark Brownlow for helping construct the ILF lighting environments, Jessi Stumpf and Andrew Jones for assembling the Caryatid model, Z-Corp for 3D printing the Caryatid, Anshul Pandey for advice regarding the HDR image assembly, Ted Chavalas of Panoscan, Inc. for lens selection advice, and Barry Lewine of Elumens, Inc. for light field reproduction discussions. We particularly thank Tomas Lochman of the Basel Skulpturhalle for providing access to the Caryatid sculpture, Roberto Scopigno and Paolo Cignoli of CNR-Pisa for providing their MeshAlign 3D mesh merging software, and Mark Ollila, Richard Lindheim, Neil Sullivan, James Blake, and Mike Macedonia for their support of this project.

This work has been sponsored by the University of Southern California, Linköping University, and U.S. Army contract number DAAD19-99-D-0046. The content of this information does not necessarily reflect the position or policy of the sponsors and no official endorsement should be inferred.


References

1. E. H. Adelson and J. R. Bergen. *Computational Models of Visual Processing*, chapter 1. MIT Press, Cambridge, Mass., 1991. The Plenoptic Function and the Elements of Early Vision.
2. Ian Ashdown. Near-field photometry: A new approach. *Journal of the Illuminating Engineering Society*, 22(1):163–180, Winter 1993.
3. J. F. Blinn. Texture and reflection in computer generated images. *Communications of the ACM*, 19(10):542–547, October 1976.
4. Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH 98*, July 1998.
5. Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. *Proceedings of SIGGRAPH 2000*, pages 145–156, July 2000.
6. Paul Debevec, Andreas Wenger, Chris Tchou, Andrew Gardner, Jamie Waese, and Tim Hawkins. A lighting reproduction approach to live-action compositing. In *Proc. SIGGRAPH 2002*, pages 547–556, August 2002.
7. Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97*, pages 369–378, August 1997.
8. Bastian Goldlücke, Marcus Magnor, and Bennett Wilburn. Hardware-accelerated dynamic light field rendering. In Günther Greiner, Heinrich Niemann, Thomas Ertl, Bernd Girod, and Hans-Peter Seidel, editors, *Proceedings Vision, Modeling and Visualization VMV 2002*, pages 455–462, Erlangen, Germany, November 2002. aka.
9. Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The Lumigraph. In *SIGGRAPH 96*, pages 43–54, 1996.
10. Gene Greger, Peter Shirley, Philip M. Hubbard, and Donald P. Greenberg. The irradiance volume. *IEEE Computer Graphics & Applications*, 18(2):32–43, March-April 1998.
11. Wolfgang Heidrich, Hendrik Lensch, Michael Cohen, and Hans-Peter Seidel. Light field techniques for reflections and refractions. In Dani Lischinski and Gregory Ward Larson, editors, *Rendering Techniques*


- '99: *Proceedings of the 10th Eurographics Workshop on Rendering (EGRW-99)*, pages 187–196, Granada, Spain, June 1999. EUROGRAPHICS, Springer.
12. Aaron Isaksen, Leonard McMillan, and Steven Gortler. Dynamically reparameterized light fields. pages 297–396, July 2000.
 13. C. Buehler L. McMillan J. C. Yang, M. Everett. A real-time distributed light field camera. In *EUROGRAPHICS Annual Conference Proceedings*, 2002.
 14. M. Koudelka, S. Magda, P. Belhumeur, and D. Kriegman. Image-based modeling and rendering of surfaces with arbitrary brdfs. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 568–575, 2001.
 15. Marc Levoy and Pat Hanrahan. Light field rendering. In *SIGGRAPH 96*, pages 31–42, 1996.
 16. Marc Levoy, Kari Pulli, Brian Curless, Szymon Rusinkiewicz, David Koller, Lucas Pereira, Matt Gintzton, Sean Anderson, James Davis, Jeremy Ginsberg, Jonathan Shade, and Duane Fulk. The digital michelangelo project: 3d scanning of large statues. *Proceedings of SIGGRAPH 2000*, pages 131–144, July 2000. ISBN 1-58113-208-5.
 17. Gene S. Miller and C. Robert Hoffman. Illumination and reflection maps: Simulated objects in simulated and real environments. In *SIGGRAPH 84 Course Notes for Advanced Computer Graphics Animation*, July 1984.
 18. I. Sato, Yoichi Sato, and Katsushi Ikeuchi. Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Transactions on Visualization and Computer Graphics*, 5(1):1–12, January-March 1999.
 19. Ph. Slusallek W. Heidrich, J. Kautz and H.-P. Seidel. Canned lightsources. In *Proceedings of the EG Rendering Workshop 1998*, 1998.
 20. Gregory J. Ward. The RADIANCE lighting simulation and rendering system. In *SIGGRAPH 94*, pages 459–472, July 1994.



(a) Real-world photograph



(b) Rendering with an ILF



(c) Rendering with a single lighting measurement


Figure 11: A scene rendered with an incident light field. (a) A real-world scene illuminated by spatially varying illumination (b) A synthetic recreation of the scene illuminated by the captured incident light field, showing consistent illumination throughout the volume of the objects. (c) The same synthetic scene rendered with a single incident illumination measurement.



(a) Real-world photograph



(b) Rendering with an ILF



(c) Rendering with a single lighting measurement

Figure 12: A second scene rendered with an ILF. (a) A real-world scene illuminated by high-frequency spatially varying illumination (b) A synthetic recreation of the scene illuminated by a 30×30 incident light field, showing lower lighting detail (c) The same synthetic scene rendered with a single incident illumination measurement.