Geometry-Corrected Light Field Rendering for Creating a Holographic Stereogram

Joel Jurik
USC ICT
Playa Vista, CA

Thomas Burnett
Zebra Imaging
Austin, TX

Michael Klug
Zebra Imaging
Austin, TX

Paul Debevec
USC ICT
Playa Vista, CA

Abstract

We present a technique to record and process a light field of an object in order to produce a printed holographic stereogram. We use a geometry correction process to maximize the depth of field and depth-dependent surface detail even when the array of viewpoints comprising the light field is coarsely sampled with respect to the angular resolution of the printed hologram. We capture the light field data of an object with a digital still camera attached to a 2D translation stage, and generate hogels (holographic elements) for printing by reprojecting the light field onto a photogrammetrically recovered model of the object and querying the relevant rays to be produced by the hologram with respect to this geometry. This results in a significantly clearer image of detail at different depths in the printed holographic stereogram.

1. Introduction and Background

Traditional holograpy involves the recording of an interference pattern of light onto holographic film which, when illuminated properly, reproduces the light field originally incident upon it during its exposure. Most commonly, a coherent reference beam formed by a defocussed laser source is split so that it illuminates both an object and, through an alternate path, the holographic film. These two coherent wavefronts produce interference patterns on the scale of the wavelength of light recorded by the high-resolution film. When the developed film is illuminated from the angle of the reference beam, it reflects the 4D light field of the object originally incident upon it. As such, a viewer can view the object from any angle and be presented with a faithful three-dimensional, autostereoscopic view of the object via the hologram. A drawback, however, is that the object appears to be illuminated by a point source of light – the original direction of the laser lighting it – rather than from a natural environment of indicent illumination. For highly polished objects, this can be a drawback, since the preferred studio lighting typically comes from a set of area light sources.

Key work in the 1990's [5, 2] presented practical techniques for capturing and rendering 4D light field data of objects using digital photography and computer graphics rendering. In [5], a camera was moved to a planar or cylindrical array of viewpoints relative to an object, and quadralinear interpolation was used to query new rays of light intersecting points (u,v) and (s,t) of two planes parameterizing the light emanating from the object. Novel viewpoints can be generated simply by querying the radiance along the set of rays comprising the pixels of any virtual camera aimed toward the object, even if it is closer or further from the object than the original array of cameras. [2] went further by acquiring light fields (or, "Lumigraphs") from unstructured arrays of viewpoints and allowing a "geometry correction" step, wherein ray queries into the scene would first be intersected with an approximate model of the geometry of the scene, and then traced back to the nearest available camera viewpoint (or more generally, viewpoints) to generate an adjusted set of rays with which to query the radiance along the ray. This had the effect of refocussing the Lumigraph onto the surface of the object, allowing for greater sharpness in the renderings. Also related is the view-dependent texture mapping technique of [1], which projected different texture maps onto an approximate photogrammetrically recovered model of a scene, yielding a depth-corrected light field when the spacing of the original camera views would become sufficiently dense.

Light field capture has received recent heightened interest with the release of the *Lytro* consumer camera, which uses a single large-aperture lens and a microlens array over a high-resolution sensor as in [6] to capture a 4D light field of a scene in a single hand-held shot. It is notable that some of the first 4D photographic data recorded for producing full-parallax holographic stereograms was recorded using a lenslet array technique [7] in the 1960's.

Digital hologram printing techniques have also advanced in step with light field photography techniques, allowing arbitrary illumination fields to be recorded onto holographic

978-1-4673-1612-5/12/$31.00 ©2012 IEEE

film. The printing technique of [4, 3] which we employ in our work uses a moving aperture plate to expose one small pixel, or *hogel*, at a time of holographic film. Each pixel measuring a millimeter or less across is exposed with a 2D image of the angularly-varying light which should radiate from that pixel. The image is projected onto a diffusing screen using a coherent laser light source, which is also split to form a reference beam also illuminating the sample. The hogels, taken as a set, comprise a light field of the object focused on the plane of the hologram. When the aperture has moved to expose all hogels, the film is developed and the holographic stereogram can be viewed when illuminated from the position of the reference beam.

The large amount of data for such digital holographic printing is most often rendered with graphics hardware using specially programmed views of texture-mapped polygon models or subjects such as machine parts, cars, buildings, or terrain. The data can also be produced using light field photography, allowing for the illumination in the scene to be anything one would desire rather than the same point source of light used to produce a reference beam. Unfortunately, there is a general mismatch between the number of photographs which can be practically acquired of an object (in our work, an array of 64×64 photographs), and the resolution of the data projected onto each hogel (in our work, 512×512 angular ray samples). As a result, producing high-quality digitally printed holograms from light field data has remained difficult.

In this paper, we show that depth-corrected light field rendering can be used to derive such high-resolution hogel data from a photographically acquired light field with a much smaller number of viewpoints. We chose a real-world object consisting of a variety of reflectance properties and lettering at different scales, with a clear need to be able to resolve details at various depths. The resulting holographic print achieves a higher level of detail and depth of field than previously achievable using the light field acquisition technique.

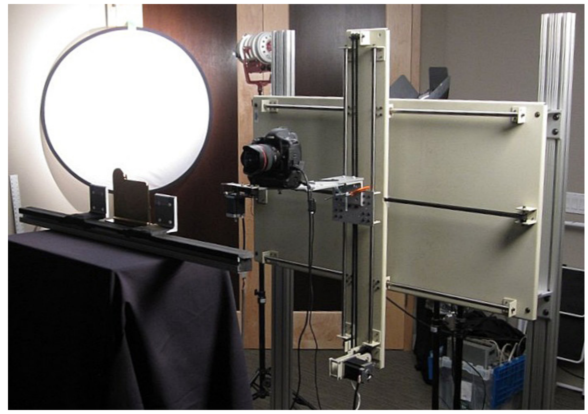

Figure 1. Light field capture setup using a high-resolution digital camera on a 2D translation stage. The object can be seen in silhouette at the bottom edge of the circular reflector.

Figure 2. One of a 64×64 array of digital photographs showing the object and fiducial markers.

2. Light Field Acquisition

We used the straightforward light field acquisition technique of attaching camera (a 16MP Canon EOS-1Ds Mark II) camera to a vertical 2D translation stage shown in Figure 1. The object we recorded – a high-relief award plaque with a protruding name plate – was chosen since it has a variety of reflectance properties including mirror-like specular, rough specular, and diffuse reflectance, and with important detail at a variety of scales and depths. We positioned the object parallel to the translation stage 43cm away and placed checker fiducial markers at the sides of the object to use in pre-processing for image stabiization. We illuminated the polished reflective surfaces using a combination of broad area light sources and ambient illumination. We used the Canon 14mm L-series lens with approximately a 104 degree by 82 degree field of view, well matched to the 90 degree by 90 degree field of view hogel data we would record onto the hologram.

We shot 4,096 images in a 64×64 grid, with a spacing of 1.135 cm between viewpoints horizontally and 1.024 cm vertically, taking approximately six hours. These camera spacings were somewhat different than we had programmed due to the weight of the camera retarding the vertical lift of the gantry, so we were fortunate to verify them through measurement. We calibrated the intrinsic parameters of our camera and lens using the checkerboard technique of [8].

3. Pre-Processing

Each of the 4,096 images are recorded in RAW format with pixel resolution 5010×3336 pixels, an example is shown in Figure 2. Uncompressed and color-interpolated, this would be 765 GB of image data, currently inefficient and impractical to manipulate. Although the object fills only a portion of the frame, there is still excess resolution for the holographic printing process which will record a 420×420 hogel image. To maximize the number of pixels on the object, the camera was placed relatively close, just making sure that it did not occlude light sources from the object during the traversal.

We undistort, crop, and rectify each image in the light field to reduce the data required. To do this efficiently, we detected the position of each of the four outer fiducials surrounding the object. To maximize the angular extents of the light field, we extrapolated fiducial positions for images in which some of the fiducials (but not all of the object) had left the frame. To reduce the light field data, we wish to crop each image to the object area between the fiducials, but we also need to eliminate the lens distortion while also minimizing image resampling. We thus compose the undistortion function with a planar homography which takes an image of the object in the original photograph and maps it into a cropped and rectified 640×640 pixel square where each of the fiducials occupy the same pixel coordinate for all images in the light field as in Figure 3. This greatly reduces the amount of data which will be necessary to be rebinned to form the hogel data for the hologram without needing to undistort each RAW image at full resolution, which would require great amounts of storage and computation. Now, the light field data takes 19GB.

Our wide-angle lens exhibited significant chromatic aberration, which produced color fringing and smeared the red channel in particular toward the corners of the image. We were able to largely eliminate the effect of the chromatic aberration by ”colorizing” each image’s relatively sharp green channel with a low-pass version of the original image’s chromaticity. Specifically, we replaced the red channel R with $G * (blur(R)/blur(G))$ and blue B with $G * (blur(B)/blur(G))$; the original green channel stayed the same. This worked since our object did not have high-frequency color variation; avoiding this for more colorful objects would require a cylindrical or spherical light field to keep the object in the center of the frame during acquisition. Finally, we sharpen the image slightly to bring out more detail.

Figure 3. 640 x 640 pixel undistorted, rectified, and sharpened images with the outer fiducials stabilized and chromatic aberration removed.

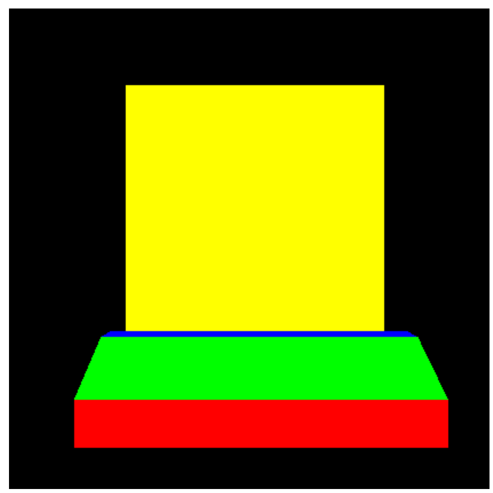
Figure 4. The photogrammetrically recovered approximate model of the object, consisting of four rectangles in 3D.

4. Hogel Reprojection

The hologram we create consists of 420×420 hogels (over about 30×30 cm), which appear as angularly varying pixels in the printed hologram. Each hogel has a resolution of 512×512 pixels covering $\pm 45°$ in both vertical and horizontal directions. So, our goal is to generate a 420×420 array of 512×512 images representing the hogels.

The straightforward way to generate the hogels is to query rays in the light field radiating from the plane of the hologram. We first choose where the plane of the hologram should be relative to our scene, which is to say where the object should appear relative to the plane of the hologram. We could naturally choose the plane of the hologram to be the plane of the fiducials, which are 43 cm in front of the light field plane from which the camera took its pictures. For each hogel position, we would query a 512×512 set of rays directed toward the hogel’s spatial position covering the $\pm 45°$ field of view of the hogel. As in [5], each ray query would be answered through quadralinear interpolation, which involves bilinear interpolation of pixel lookups into the four nearest cameras to the queried ray. Unfortunately, this does not produce an especially sharp rendition of object details away from the plane of the hologram, as seen in Figure 6.

In order to create a hologram of the object that is as sharp as possible, we must refocus these hogels. To do this we must first approximate the geometry of the object. We decided to choose four planes of the object as seen in Figure 5. We measured the four corners of each plane by using photogrammetry. This entails placing two virtual cameras at known distances and projecting rays into the world and finding their 3D intersection at the plane corners in the image. Once we recover the world coordinates of the corners of all planes of the object, we choose the placement of the focal plane. We decided to place it parallel to the camera plane, and in the middle of the object plane with the protruding block of text. Looking at Figure **??**, for each hogel in the resulting hologram, we place a virtual camera on the focal plane facing the camera plane, point P on the diagram. Then we project rays into the world from the virtual camera and find the intersection point Q with the geometry of the object that we calculated. We then find the intersection point C with the camera plane. Since we only have 64x64 camera positions, there is a strong likelihood that the intersection with the camera plane will not fall exactly on one of the cameras, therefore we must find the four closest cameras. For each camera, we find the pixel using bilinear interpolation that corresponds to the intersection point Q of the geometry of the object, then we follow the orange arrows in the diagram blending the pixels to better approximate the radiance at P toward C. We do this 512x512 times for each virtual camera. To eliminate some of the background and the fiducials, we subtract the average pixel value of the black blanket from each pixel, and then only include parts of the image that are close to the geometry of the object.

5. Memory Management

This project uses 4096 images with resolution of 640x640. We choose a 640x640 image since it is close to the 420x420 spatial resolution of the hologram, and it is not too large as to make memory management impractical. To be as precise as possible, we use the pfm image format. The pfm format is similar to the ubiquitous ppm format, except that it contains float values for each pixel, requiring 4 bytes for each of the red, blue, and green channels. This totals to roughly 19 gigabytes of data. Since storing 19 gigabytes of data in memory is not always feasible, we came up with a least recently used (LRU) scheme to slightly speed up the operations. This is useful because many of the operations may use the same image already stored in memory, and iterating through the program in a certain way will help exploit that fact. Light fields are also an excellent candidate for a multi-threaded application, as most of the work can be done independently. Together with memory management and using multi-threads, we were able to create data for the hologram in a timely fashion.

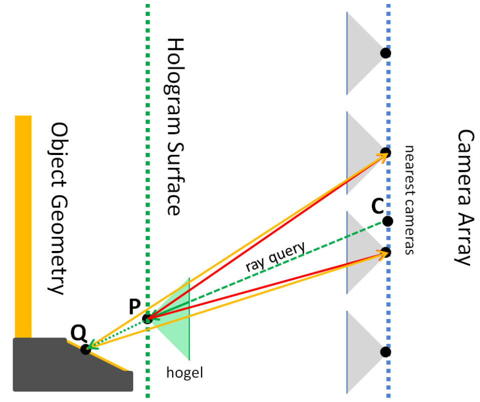

Figure 5. Our goal is to determine the radiance that the hogel at point P should exhibit along ray PC. We could naively intersect PC with the camera array plane to determine the nearest cameras and then take a weighted blend of the red rays from point P on the hologram to these cameras. Instead, we can achieve a better-focussed result by intersecting PC with the object geometry to find Q. We then take a blend between pixels from the nearest cameras aimed toward point Q (orange rays) rather than point P to better approximate the radiance at P toward C.

6. Results

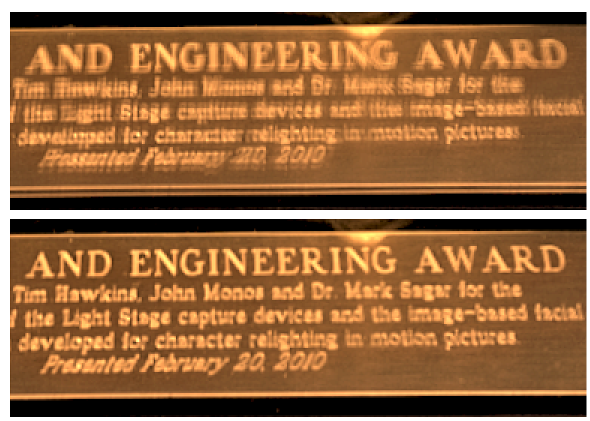

Figure 6. Top: Using a single focal plane through the main vertical surface of the object produces blurred surface details. Bottom: Using a geometry-corrected focal surface for reprojecting the hogels allows surface details to remain sharp.

Figure 6 shows the result of a synthetically rendered hologram both with and without the geometric depth correction step. The top image shows a rendering made by reprojecting the light field data without using the object's geometry, instead refocussing the image data onto a flat plane coindicent with the main plaque part of the object. As a result, the limited resolution of the camera array becomes evident by blurry and aliased image details, making the text difficult to read. The bottom image of Fig. 6 shows the result of using the object geometry to refocus the light field data, which allows the sparsely sampled angular measurements to line up with each other and produce sharp text across the entire surface of the steeply sloped plane.

We finally used the holographic printing process of [4] to produce real printed holographic stereograms of the object using geometry correction. First, we created 420×420 hogels with 256×256 angular resolution, with the main vertical plaque of the object coincident with the hologram surface. In the printed 30 cm $\times$ 30 cm result, the protruding text was mostly legible, but there was some blurring due to the limited hogel resolution. We then generated a second hogel dataset of 420×420 hogels with 512×512 angular resolution, and we moved the hologram surface to be coincident with the front of the protruding base of the object. Two photographs of this hologram (illuminated from the proper illumination direction) can be seen in Figure 7. The small text on the sloped plane of the base is sharp and legible, which would not have been possible without the geometry-correction reprojection step. The main vertical plaque is legible with its larger text, but still slightly blurry. This is because even with the higher hogel resolution, the background was still too far behind the plane of the hologram for the holography process to produce a sharp result with a finite-sized light source. We will thus be printing one additional hologram with the plane of the hologram through the middle of the object, which we expect to be entirely sharp for all surfaces.

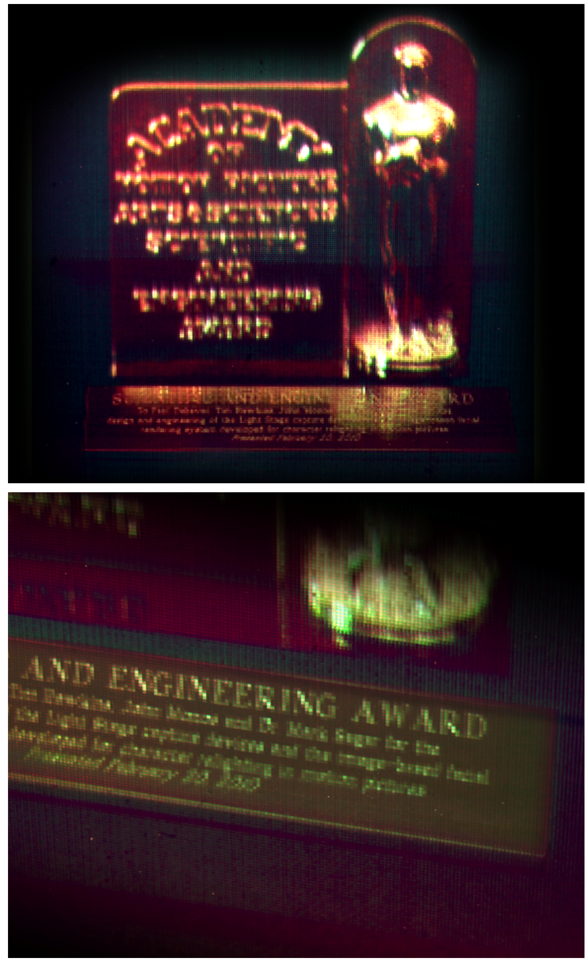

Figure 7. The hologram.

7. Conclusion and Future Work

In this work, we have shown a practical way of recording and processing light field data used to create a high-quality holographic print showing arbitrary illumination conditions in the scene. By using geometry correction, we can refocus a light field acquired from far fewer viewpoints than the angular resolution of the hogels would appear to require. The resulting hologram has sufficient image detail to observe high-resolution details of an oject at a variety of depths.

Acknowledgements

This work was sponsored by the University of Southern California Office of the Provost and the USC Provost's Ph.D. Fellowship. The content of the information does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

References

[1] P. E. Debevec, C. J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, SIGGRAPH '96, pages 11–20, New York, NY, USA, 1996. ACM. 1

[2] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, pages 43–54, Aug. 1996. 1

[3] M. Klug. Scalable digital holographic displays. In *Proc. PICS 2001: Image Processing, Image Quality, Image Capture Systems Conference*, April 2001. 2

[4] M. A. Klug, M. Holzbach, and A. Ferdman. US Patent 6,330,088: Method and apparatus for recording one-step, full-color, full-parallax, holographic stereograms, Dec 2001. 2, 5

[5] M. Levoy and P. M. Hanrahan. Light field rendering. In *Proceedings of ACM SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, pages 31–42, 1996. 1, 3

[6] R. Ng. Fourier slice photography. *ACM Trans. Graph.*, 24(3):735–744, July 2005. 1

[7] R. V. Pole. 3-d imagery and holograms of objects illuminated in white light. *Applied Physics Letters*, 10(3):20–22, 1967. 1

[8] Z. Zhang. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(11):1330–1334, Nov. 2000. 2