Calibration of a Surface Light Field Capture System

Jan-Michael Frahm, Greg Coombe and Anselmo Lastra
University of North Carolina at Chapel Hill

Abstract

In this paper we introduce a simple procedure to calibrate a surface light field capture system. This system was developed to capture the illumination from a virtual environment map [2]. However, the system has the drawback of a complex calibration procedure that is limited to planar screens. We propose a simple calibration procedure using a reflective calibration object that is able to deal with arbitrary screen geometries. Our calibration procedure is not limited to our application and can be used to calibrate most camera projector systems.

1. Introduction

Surface light fields (SLFs) [6, 4] are used to represent the exitant radiance of an object under fixed illumination. To construct an SLF, images are captured from a set of viewpoints and projected onto the surface of a known geometric model. This parameterization results in a compact representation that enables the capture and display of complex, view-dependent illumination of real-world objects. Surface light fields can be easily captured with minimal equipment [3], compressed [5], and interactively rendered [1].

One limitation of SLFs is that the desired illumination must be physically recreated during the capture process. This is problematic for synthetic environments such as games or virtual environments. Our solution is to illuminate the real object with a high dynamic range environment map using a projector-camera system [2]. The system consists of a projector that casts synthetic illumination from a light probe onto a screen. The light is reflected from the screen onto the object and captured by a camera. This simulates the object being lit by the synthetic environment as can be seen in Figure 1. The system is shown in Figure 2.

However, the proposed system [2] has the drawback of a complex calibration procedure which is limited to planar screens. The components that have to be calibrated are the position and the geometry of the screen in the coordinate system of the tracking board. We propose a simple geometric calibration procedure that uses a reflective calibration object and is able to deal with arbitrary screen geometries.

Figure 1. Left: pitcher model in St. Peter's. Right: heart model in Uffizi. (courtesy debevec.org)

2. Calibration of the projector camera system

The goal of the calibration procedure is to determine the light rays corresponding to each projector pixel. These rays are used to index into the environment map to provide a "window" into the synthetic environment. This relationship is influenced by the angle between object and screen, geometry of the screen, and nonlinearities of the projection. These factors imply that physically measuring the system is difficult and often inaccurate. We propose a calibration procedure that takes advantage of a particular property of light probes; they are independent of translation. Thus the calibration procedure needs to only compute the rays emanating from the object, and does not need to compute the translation information. This simplifies the calibration procedure to determining the relation between a pixel on the screen and a ray in object space.

Light Ray Calibration We place a mirrored sphere with known radius and position in approximately in the same location as the object. This mirrored sphere reflects the illuminated points on the screen back to the camera. Using the known projector pixel and the reflection of this point into the camera, the ray associated with each projector pixel can be computed. This is shown in Figure 3.

The process works as follows. A pixel is illuminated on

Pixel	-2%	-1%	+1%	+2%
(34, 131)	0.4064°	0.2011°	0.1971°	0.3902°
(151, 63)	2.0260°	0.9988°	0.9719°	1.9179°
(-18, 156)	0.0919°	0.0455°	0.0446°	0.0883°
(-72, 154)	0.7546°	0.3733°	0.3655°	0.7235°

Table 1. Angular error vs. relative radius error

the screen, and an image is captured of the light reflected from the sphere. The position of the pixel in the image is segmented using an adaptive threshold. The ray through the camera plane into the scene can be computed from the pose of the tracking board. This ray is then traced into the scene and intersected with the sphere. At the intersection point, the normal is computed and a reflected ray is generated. This reflected ray is the ray in 3D space which corresponds to the projector pixel.

This establishes a correspondence between projector pixels and rays in the scene. For rendering, these rays are rotated according to the delta rotation of the tracking board, and used to index into the environment map. The resulting images for the projector are shown in Figure 3.

Error The primary source of error in this calibration procedure is from the physical measurement of the sphere. We measured the error in the reflected rays as a function of the error in radius of the sphere for several projector pixels. Pixels near the edge of the sphere are more susceptible to error, since the normals change more quickly. Minor sources of error which we did not measure are camera pose estimation and segmentation problems.

3 Future Work

Currently, we use a 9x9 sampling of the screen (100 calibration points) aligned on a regular grid. For screens which are mostly planar, this is sufficient. However, for more complex screens a denser sampling would be needed. It would be interesting to allow the user to define specific screen features by selecting projector pixels, and tessellating the screen geometry. It might also be possible to develop a fully-automatic adaptive system which illuminates points in the environment and measures how far off they are from the predicted point in the camera image.

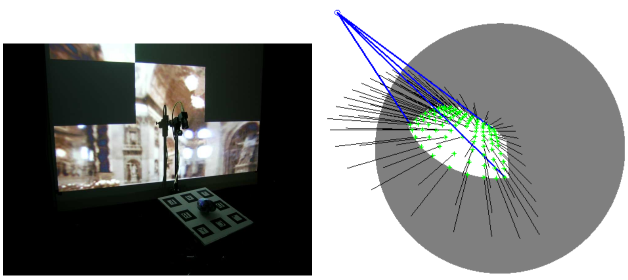

Figure 2. Left: The surface light field capture system. Right: Reflected light rays from calibration object.

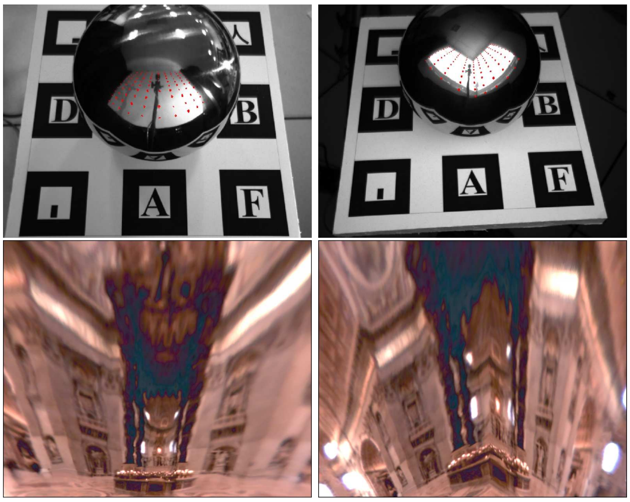

Figure 3. Top: Reflection of the screen on the calibration object (planar on the right, corner on the left) with segmented pixel positions. Bottom: Light probe mapped onto the screen (planar on the right, corner on the left).

References

[1] W.-C. Chen, J.-Y. Bouguet, M. Chu, and R. Grzeszczuk. Light field mapping: Efficient representation and hardware rendering of surface light fields. In *SIGGRAPH*, 2002.

[2] G. Coombe, J.-M. Frahm, and A. Lastra. Capturing Surface Light Fields of Real Objects Under Virtual Illumination. Technical report TR06-001, University of North Carolina at Chapel Hill, May 2006.

[3] G. Coombe, C. Hantak, A. Lastra, and R. Grzeszczuk. Online construction of surface light fields. In *Eurographics Symposium on Rendering*, 2005.

[4] G. S. P. Miller, S. M. Rubin, and D. Ponceleon. Lazy decompression of surface light fields for precomputed global illumination. In *Eurographics Workshop on Rendering*, 1998.

[5] G. Mueller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein. Acquisition, synthesis and rendering of bidirectional texture functions. In *Eurographics*. State of the Art Reports, 2004.

[6] D. Wood, D. Azuma, W. Aldinger, B. Curless, T. Duchamp, D. Salesin, and W. Stuetzle. Surface light fields for 3d photography. In *SIGGRAPH*, 2000.