

A Photometric Approach to Digitizing Cultural Artifacts

Tim Hawkins

Jonathan Cohen

Paul Debevec

University of Southern California Institute for Creative Technologies¹

ABSTRACT

In this paper we present a photometry-based approach to the digital documentation of cultural artifacts. Rather than representing an artifact as a geometric model with spatially varying reflectance properties, we instead propose directly representing the artifact in terms of its *reflectance field* – the manner in which it transforms light into images. The principal device employed in our technique is a computer-controlled lighting apparatus which quickly illuminates an artifact from an exhaustive set of incident illumination directions and a set of digital video cameras which record the artifact's appearance under these forms of illumination. From this database of recorded images, we compute linear combinations of the captured images to synthetically illuminate the object under arbitrary forms of complex incident illumination, correctly capturing the effects of specular reflection, subsurface scattering, self-shadowing, mutual illumination, and complex BRDF's often present in cultural artifacts. We also describe a computer application that allows users to realistically and interactively relight digitized artifacts.

Categories and subject descriptors: I.2.10 [Artificial Intelligence]: Vision and Scene Understanding - *intensity, color, photometry and thresholding*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - *color, shading, shadowing, and texture*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism - *radiosity*; I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture - *radiometry, reflectance, scanning*; I.4.8 [Image Processing]: Scene Analysis - *photometry, range data, sensor fusion*. **Additional Key Words and Phrases:** image-based modeling, rendering, and lighting.

1 Introduction

Creating realistic computer models of cultural artifacts can aid their study by remote parties as well as serve as an improved record of the artifact for archival purposes. While a photograph faithfully records an artifact's appearance from a single point of view in a particular lighting environment, it can be far more informative to be able to see the artifact from any angle and in any form of lighting - allowing a scholar to give it greater scrutiny as well as the general observer to see the artifact in its natural environment.

The standard computer model of a cultural artifact consists of a geometric surface model covered by a texture map which represents the artifact's spatially varying diffuse reflectance properties. Sometimes, a specular component is added to render the shininess, often set to a single representative value for the object. An artifact's geometry is most commonly acquired using active sensing such as laser scanning or structured light, and the texture maps are usually constructed using a simplified form of reflectometry: lighting the object from a known direction with a calibrated light source, and then, using an estimate the object's surface orientation, determining the spatially varying albedo of the object. This basic approach has produced numerous high-quality records and render-

ings of cultural artifacts such as statues, vases, and interior environments.

Unfortunately, this traditional approach is difficult to apply to a large class of cultural artifacts – ones that exhibit complex reflectance properties such as anisotropy or iridescence, ones that exhibit significant self-shadowing or mutual illumination, ones that exhibit significant subsurface reflection, objects that are highly specular or translucent, and objects with intricate surface geometry. We can illustrate these difficulties with several examples:

- A fur headband would be difficult to digitize since it does not have a well-defined surface – the stripe of a laser scanner or video projector would scatter through the fur rather than drape over it, and the surface reconstruction algorithm would have difficulty reconstructing a surface from the data. Even if the fur headband's geometry could be captured, it could be impractical to represent the geometry of tens of thousands of individual hair segments, and to subsequently realistically render the effects of this fur being illuminated. Similar challenges exist for digitizing many types of clothing as well as for human hair.
- A small jade sculpture would exhibit significant subsurface scattering - light hitting it from behind would cause it light up from the front, and light striking the front would penetrate the surface a considerable distance. This could complicate the range scanning process, but would pose an even more significant problem for reflectometry: current techniques neither estimate nor represent subsurface scattering properties of an object. A computer rendering of the jade sculpture might more closely resemble painted green rock than luminous jade.
- An intricately carved ivory sculpture with complex internal geometry would be challenging to digitize due to both significant self-shadowing and mutual illumination. Scanning its geometry would be complicated by the self-shadowing; narrow crevices are difficult to record with triangulation-based scanning methods in which each surface point must be visible to both the source of light and the image sensor. Furthermore, the complexity of the geometry could complicate the range map merging process, which works best when there are relatively coherent surface sections. Mutual illumination – the fact that light will bounce between surfaces inside the object's reflective concavities – would complicate the reflectometry stage since it becomes very difficult to reliably control the illumination incident on any particular surface. The significant mutual illumination will also complicate rendering, requiring expensive global illumination algorithms to produce images that correctly replicate the appearance of the original object.
- A polished silver necklace, encrusted with rubies, diamonds, and abalone, would also be difficult to digitize. The low diffuse reflectance of the silver and the internal reflections and refractions in the rubies and diamonds would make geometry capture with structured light or laser scanning impractical. It would be hard to measure the reflectance of the reflective surfaces since most reflectometry techniques are better suited to materials with a significant diffuse reflection component. The translucence of the rubies and diamonds would make modeling their reflection characteristics difficult, and the reflectance

¹USC Institute for Creative Technologies, 13274 Fiji Way 5th floor, Marina del Rey, CA, 90292. Email: {timh, debevec}@ict.usc.edu.

of the iridescent abalone would be too complex to represent with most currently available reflectance models.

In summary, a large class of cultural artifacts exhibiting complex geometric and reflectance properties cannot be effectively digitized using currently available techniques. This poses a significant problem for the application of computer graphics to cultural heritage, as many of the materials and designs used by craftspeople and artisans are specifically chosen to have complex geometry or to reflect light in interesting ways.

In this paper, we show that an alternative approach based on capturing *reflectance fields* [4] of cultural artifacts can acquire, represent, and render any of the above artifacts just as easily as it could a clay jug or granite statue, with simple acquisition and rendering processes, and can produce photorealistic results. The technique is data intensive, requiring thousands of photographs of the artifact, and as currently applied allows only for relatively low-resolution renderings. As such we discuss its advantages and disadvantages over current techniques as well as potential hybrid methods.

In our proposed technique the artifact is photographed under a dense sampling of incident illumination directions from a dense array of camera viewpoints. The device we use to acquire this dataset is a *light stage* consisting of a semicircular rotating arm with an array of strobe lights capable of illuminating an object placed at its center from up to a thousand different directions covering the entire sphere of incident illumination. Images taken with digital cameras are compressed together into a single multi-dimensional dataset comprising the object's *reflectance field* [4], which characterizes how the object transforms incident illumination into radiant imagery. Renderings of the object can then be created under any form of illumination - such as the light in a forest, a cathedral, a museum, or in a candlelit hut, by taking linear combinations of the images in the reflectance field dataset. No geometric model of the object is required, and the resulting renderings capture the full complexity of the object's interaction with light, including self-shadowing, mutual illumination, subsurface scattering, and translucency, as well as non-lambertian diffuse and anisotropic specular reflection.

In this paper we apply this process to a number of Native American cultural artifacts including an otter fur headband, a feathered headdress, an animal-skin drum, and several pieces of neckwear and clothing. We show these artifacts illuminated by several real-world natural lighting environments, and describe a software program for interactively re-illuminating artifacts in real time. We also describe the reflectance field capture process and the equipment involved. We also propose using view interpolation to extrapolate a discrete set of original viewpoints to of an artifact to arbitrary novel viewpoints in conjunction with the reflectance field capture technique.

The central contribution of this paper is to demonstrate that the reflectance field capture technique, introduced in the context of rendering human faces in [4], provides particular advantages over current digitization techniques when applied to cultural artifacts. Since artifacts can have stronger specular components than human skin, we also show that capturing shiny artifacts requires the acquisition of high dynamic range [5] reflectance field image data. We present an interactive program for visualizing re-illuminated reflectance fields, and suggest how view interpolation could be used to continuously vary the viewpoint of an artifact from a discretely sampled set of viewpoints. Finally, we suggest improvements to the light stage apparatus specifically for acquiring complete view-independent models of cultural artifacts.

2 Background and Related Work

Current leading techniques for digitizing three-dimensional cultural artifacts involve acquiring multiple range scans of the artifact, as-

sembling them into a complete geometric surface model, and then using a form of reflectometry to derive lighting-independent texture maps for the artifact's surfaces.

Range scans are acquired most commonly through laser-stripe scanning, by projecting patterns of light from a video projector, or through illumination-assisted stereo correspondence. Individual range scans can be aligned to each other using Iterated Closest Points (ICP) algorithm [23] and merged using either polygon zipping [19] or volumetric range merging [2]. A sampling of recent projects which have used these techniques to derive geometric models of cultural artifacts are IBM Watson's Florentine Pieta Project [17], Stanford's Digital Michelangelo Project [11], Electricit de France's Cosquer cave (1994) and Colossus of Ptolemy (1997) projects, the Canadian National Research Council's museum artifact scanning work [1], and work to scan vases [16] at the Istituto di Elaborazione dell Informazione in Pisa.

Reconstructing an artifact's appearance - not just its geometry - is the second stage of the digitizing process. Some techniques (e.g. [6], [22]) directly project photographs of the object onto the artifact's geometry to form diffuse texture maps; this technique has the advantage that the renderings will exhibit realistic, pre-rendered shading effects including self-shadowing and mutual illumination, but have the disadvantage that the lighting is fixed according to the conditions in the original imagery. For environments it is sometimes acceptable to have static lighting; for artifacts static lighting is generally less acceptable since the directions of incident illumination on an artifact change as the artifact is rotated and there is often a desire to visualize an artifact as it would be seen in different illumination environments.

Deriving lighting-independent texture maps for artifacts has been done in several projects. One set of techniques lights an object from one or more directions and uses the geometric model to estimate the surface's diffuse albedo for all points on the surface. [13] solved for the spatially varying diffuse reflectance properties across a diffuse object using different lighting and observation directions. Related techniques are used in [17] and [11] to derive illumination-independent texture maps for marble statues. [1] uses the intensity return of its collimated tri-colored laser scanner to derive color lighting-independent texture maps. Recovering specular object properties has been investigated in [18].

These current techniques have produced excellent 3D models of artifacts with well-defined surface geometry and generally diffuse reflectance characteristics. However, these current techniques for model acquisition are difficult or impossible to apply for a large class of artifacts exhibiting complex surface microstructure, spatially varying specular reflection, complex BRDFs, translucency, and subsurface scattering. As a result, artifacts featuring silver, gold, glass, fur, cloth, jewels, jade, leaves, or feathers are very challenging to accurately digitize and to convincingly render.

Recent work [4] building upon related image-based rendering techniques [9, 15, 21] presented a technique for creating relightable computer models of human faces without explicitly modeling their geometric or reflectance properties. Instead, the face was illuminated from a dense array of incident illumination and a set of digital images were captured to represent the face's reflectance field. The images from the face's reflectance field were then combined together in order to produce images of the face under any form of illumination, including lighting environments captured from the real world as in [3]. [12] applied this technique in rendering cultural artifacts exhibiting diffuse reflectance properties. In this paper, we show that this technique can be applied to the digitization of cultural artifacts exhibiting any geometric or reflectance properties including those that are traditionally difficult to model and render. We discuss issues which arise in applying these techniques to capturing cultural artifacts, including the need to acquire image data sets using high dynamic range photography

[5] to properly capture and render specular highlights. We furthermore describe an interactive lighting tool that allows artifacts to be re-illuminated by a user in real time, and propose image-based rendering techniques that will allow an artifact to be manipulated in 3D as well as being arbitrarily illuminated. In this work we use a collection of Native American clothing and jewelry to demonstrate the possibilities of the technique.

3 Dataset Acquisition



Figure 1: **The Light Stage** illuminates an artifact (center) from a dense array of incident illumination directions as its appearance is recorded by high-speed digital video cameras (one can be seen in the upper left). This quarter-second photographic exposure shows several of the lights on at once although only one strobe light flashes at any given time.

The data acquisition apparatus used in this work consists of a semicircular arm three meters in diameter that rotates about a vertical axis through its endpoints. Attached to the arm are twenty-seven evenly spaced xenon strobe lights, which fire sequentially at up to 200 Hz as the arm rotates around the subject. The arm position and strobe lights are computer-controlled allowing the strobes to synchronize with high-speed video cameras. We have used two models of high-speed video cameras in this work. The first is the UniQ Vision UC-610; it has a single image sensor of 660 by 494 pixels and can run asynchronously at up to 110 frames per second, producing digital output to the computer. The second camera is a Sony DXC-9000, which has separate 640×480 image sensors for red, green, and blue channels and which runs progressively at 60 frames per second. The Sony camera produces sharper images and more vibrant colors, although it has analog video output which must be digitized by the computer.

Figure 1 shows the light stage capturing the reflectance field of a Lakota Native American headdress; the capture process takes approximately fifteen seconds to acquire 1,728 images of the artifact. The 1,728 images are arranged as an array of 64 directions of longitude corresponding to the rotation of the arm and 27 directions of latitude corresponding to the individual strobe lights. An alternative low-cost light stage apparatus presented in [4] consists of a single traditional light source on a two-axis rotation mechanism that spirals around the artifact, beginning at the north pole and spiraling down along the surface of a sphere to the south. This alternative device is less expensive to construct, but the device we present in this current paper allows the lights to be positioned with greater precision and repeatability and the datasets to be acquired significantly more rapidly, and with much less work.

Figure 2 shows a captured reflectance field dataset of the headdress. Images near the top of the figure are illuminated from above; images in the center of the figure are illuminated from straight forward, and to the right of the center are illuminated from the right. Images at the far left and right of the figure are actually illuminated from various directions *behind* the artifact; these images are important to capture since they show light grazes along the sides of the artifact and shines through the artifact’s translucent areas. The actual dataset taken by the apparatus is considerably higher resolution; the figure shows just every fourth image in both the horizontal and vertical dimensions.

For artifacts exhibiting relatively shiny reflectance, it becomes necessary to record the reflectance field dataset in *high dynamic range* – with imagery that can capture a greater ratio between light and dark areas than single exposures from the video cameras. For this we can employ a high dynamic range image acquisition method as in [5] to combine both under- and over-exposed images taken from the same viewpoint in the same lighting into images that capture the full dynamic range of the artifact under that lighting. To take these images, we capture the reflectance field dataset more than once; we first record it normally and then next time place neutral density filters over the video camera lenses in order to reduce the exposure of the imagery. In these subsequent passes, the images are darkened to the point that the specular highlights will be properly imaged without saturating the image sensor. In this work, we use 3-stop neutral density filters which reduce the exposure by a factor of eight. If the specular highlights still saturate the image, the exposure can be further reduced by adding additional neutral density filters or setting the cameras to a smaller aperture. In Section 5 we show an example of using this procedure to faithfully reproduce specular reflections in a synthetically illuminated artifact.

4 Illuminating Reflectance Fields of Artifacts

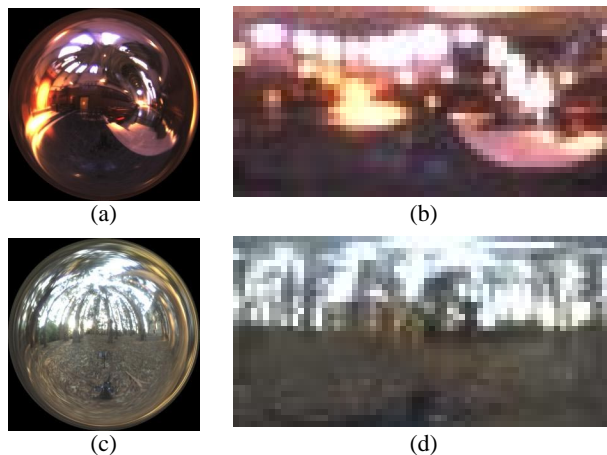


Figure 3: **Real-World Lighting Environments** (a) A light probe [3] image records the illumination in San Francisco’s Grace Cathedral (b) The probe image is resampled into a latitude-longitude format having the same coordinate system and resolution as the captured reflectance field in Figure 2. (c) A light probe image recording the incident illumination in a Eucalyptus grove. (d) The resampled version of (c). The incident illumination datasets were recorded by taking omnidirectional high dynamic range images of real lighting environments.

The motivation for our approach is to be able to realistically show digitized artifacts illuminated by any desired form of illumination. This ability can allow scholars to study how an artifact responds to light and how it may have appeared to people of the corre-




Figure 2: **A Reflectance Field Dataset** This mosaic of images of the headdress shows a sampling of the 1,728 images acquired in a light stage capture session (the original 64×27 dataset is shown as a 16×8 dataset.) The dataset shows the headdress illuminated from all possible directions of incident illumination.

sponding culture in the artifact’s natural environments. It can also allow the artifact to be realistically integrated into virtual cultural recreations or virtual museums, illuminating it with the specific illumination present in any given virtual environment.

To re-illuminate the artifacts, we employ the reflectance field illumination process originally applied for relighting human faces in [4]. First, an image of an incident illumination environment is captured or rendered; for this, the light probe technique presented in [3] can be employed. In this technique, a series of differently exposed images of a mirrored ball are combined to produce a high dynamic range omnidirectional image that measures the color and intensity of the illumination arriving from every direction in the environment. The two light probe images used in this paper are shown in Figure 3.

To create a rendering of an artifact as it would appear in such a sampled lighting environment, the light probe image is resampled into the same coordinate space and resolution as the artifact’s reflectance field; in our work this is a 64×27 image in a latitude - longitude coordinate system. The two images on the right of Figure 3 show the Grace Cathedral and Eucalyptus Grove lighting environments resampled into this coordinate system and resolution.

The next step is to multiply the red, green, and blue color channels of each of the reflectance field dataset images by the red, green, and blue colors of the corresponding pixel of the resampled lighting environment. For example, suppose that the pixel in the lighting environment corresponding to light coming directly from the right of the artifact is bright yellow. Then the reflectance field image illuminated from this direction in the light stage will be scaled so that it too is correspondingly bright and yellow. Thus, each image in the dataset becomes an accurate rendering of how the artifact would appear if illuminated by just its corresponding direction of light in the environment. The illuminated reflectance field dataset for the headdress is shown in Fig. 4.

The final step is to sum all of the images in the illuminated reflectance field dataset, producing a final rendered image showing the artifact as it would appear as illuminated by the entire sampled lighting environment at once. This procedure works because of the additive nature of light [9]: if an artifact is illuminated by

two sources of light and photographed separately as illuminated by each, then the sum of these two images will show what the artifact will look like as if illuminated by both sources at once. This assumes that the cameras taking these images are radiometrically calibrated; i.e. that the pixel values in each image are proportional to the amount of light received by the image sensor; we use the method of [5] in order to perform this calibration.

Figure 6 shows the result of illuminating the reflectance field dataset of the headdress by the two lighting environments in Figure 3 as well as a user-constructed lighting environment made with the interactive relighting tool described in Figure 5.

We note that this illumination technique yields results quite close to the physically correct answer of how the artifact would appear in the given light. Since the rendered image is a linear combination of real images of the artifact, the rendering will exhibit correct real-world illumination effects including those of anisotropic specularities, iridescence, self-shadowing, mutual illumination, subsurface scattering, translucency, and complex surface microstructure; and thus is a faithful rendering of what the artifact would look like in the specified environments. This is notable given that there is no modeling of the artifact’s surface geometry nor any steps to derive reflectance data for the artifact.

It should be mentioned that the technique will not yield perfect results in all cases. Scenes with very concentrated light sources should produce very sharp shadows on the artifact; such shadows will become slightly blurred using this technique since the reflectance fields are acquired at a finite resolution. The technique may also produce incorrect results if the spectrum of either the illumination or of the artifact’s reflectance has significant spectral peaks or valleys; carrying out the calculation solely on trichromatic RGB pixel values may fail to yield precise color balance in the renderings in such cases. Finally, this technique assumes that the artifact is illuminated by an even field of illumination; additional data acquisition and rendering would be required to show an artifact as it would be illuminated by dappled light or in partial shadow, or in close proximity to other artifacts or light sources¹.

¹This additional data could be recorded by using pixel-addressable video




Figure 4: **Lighting a Reflectance Field Dataset** A reflectance field dataset is illuminated by coloring each image in the dataset according to the color and intensity of the illumination coming from the corresponding direction in a sampled lighting environment. The images in this figure are colored according to the illumination captured in Grace Cathedral (see Figure 3 using the light probe technique in [3]). A final image of the artifact as illuminated by the environment is obtained by adding together all of the transformed images in the dataset; such an image can be seen in Figure 6.

4.1 Interactive Reflectance Field Illumination

We have written an interactive computer program that implements our relighting technique in real time on a standard PC equipped with an OpenGL graphics card. Screen snapshots from this program are shown in Figure 5. The program operates in two modes; in the first, the user can choose from a variety of captured lighting environments with which to illuminate the artifact, and can rotate the environment about the y-axis to see the light from the environment reflect differently off of the artifact. In the second mode, the user can construct the lighting environment by hand using a number of light sources. For each source, the user selects its intensity, color, and direction, as well as whether the light is a hard source such as a point light or a soft source such as an area light. As the user moves the lights, the appearance of the artifacts updates interactively at over twenty frames per second on contemporary mid-range PCs. The demo uses compressed versions of the reflectance field datasets in order to compute the renderings; this degrades the quality of the renderings slightly but makes it possible to process the quantity of data necessary at interactive rates. We have found that being able to interactively control realistic illumination significantly helps a user sense both the geometric and material properties of an artifact.

In the work so far we only render the artifact from a static viewpoint. In the next section we describe how the artifact can be rendered from novel viewpoints in addition to being rendered under arbitrary illumination.

5 Results

We have applied the photometric digitization technique to a variety of cultural artifacts chosen to exhibit complex geometry and reflectance properties that would be difficult to capture using traditional digitization techniques. Some descriptive information on the artifacts themselves is presented in Section 8.

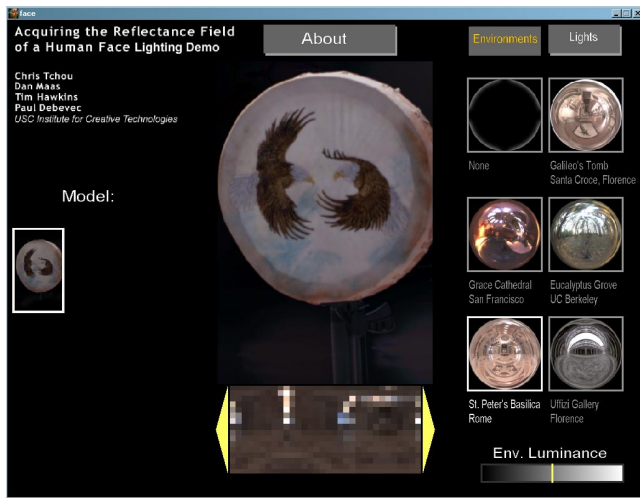
projectors, rather than uniform light sources, to illuminate the artifact.

Figure 6 shows a Lakota headdress synthetically illuminated by light captured from the Grace cathedral and Eucalyptus grove environments in Figure 3, as well as a user-specified lighting environment. Despite the headdress' complex and in places fuzzy geometry and its generally complex reflectance properties, the headdress appears realistically illuminated by each environment. A rendering of the headdress from a novel viewing angle is shown in Figure 10.


Figure 8 shows an otter skin cap and an animal bone choker necklace being captured and then illuminated by the Grace cathedral environment. The figure illustrates the need to capture the reflectance fields of shiny artifacts in high dynamic range - using multiple varying exposures to properly record both albedo reflection and specular highlights. Image (c) was illuminated using only standard single-exposure imagery in which highlight values were clipped, and yields a diffuse appearance to the reflective abalone necklace decoration. Image (d) used high dynamic range images to illuminate the artifact and thus properly replicates the shininess of the abalone.

Figure 7 shows two original light stage images and two synthetically illuminated renderings of a flat drum. Since the skin of the drum is thin and translucent, it becomes significantly illuminated even when lit from behind. Since this effect is captured in the reflectance field, the drum head will properly light up when illuminated from behind by a sampled illumination environment such as by the bright yellowish altar in the Grace cathedral lighting environment.

Figure 9 shows an approach to photometrically recording jewelry and clothing by having them be worn by an actual person. The light stage is outfitted with a chair and the reflectance field of the subject wearing the clothing is captured. Capturing clothing and jewelry in this manner can help us to better understand its physical and artistic design, and greater context is provided when the clothing is worn by a member of the culture from which the artifacts originate.



(a)



(b)

Figure 5: **An Interactive Illumination Program** The program seen above allows reflectance fields of cultural artifacts to be interactively illuminated with (a) different sampled illumination environments or (b) arbitrary user-specified illumination.



(a)



(b)



(c)



(d)

Figure 6: **Synthetically Illuminating an Artifact** (a) One of 1,728 original images in the reflectance field dataset of a headdress. (b) The headdress synthetically illuminated by the environmental lighting captured in Grace cathedral in Figure 3. (c) The headdress synthetically illuminated by the environmental lighting captured in a eucalyptus grove in Figure 3. (d) The headdress synthetically illuminated by a user-constructed lighting environment using the interactive relighting program seen in Figure 5. The renderings exhibit all of the artifact's complex properties of specularly, anisotropic reflection, translucency, and mutual illumination; such effects are usually challenging to model, represent, and render using currently available techniques for artifact digitization.




Figure 7: Capturing Translucency (a) The reflectance field of a Flat Drum is captured in the light stage. (b) When the drum is lit from behind, the front of the drum lights up due to the translucency of the drum head. The shadow of the strings that tighten the drum is also visible. (c) This image of the reflectance field of the drum illuminated by the Grace cathedral lighting environment in Figure 3 does not reveal the translucency since the environment is situated so that more light strikes the front of the drum than the back. (d) Rotating the lighting environment so that more of the light comes from behind the drum reveals the drum's translucency. A traditionally scanned 3D model of the drum with a diffuse texture map would not be able to reproduce this effect.




Figure 8: The Need for High Dynamic Range Datasets (a) a dataset image of an otter headband and an abalone shell necklace. The bright reflection in the necklace is too bright to be captured accurately by the video camera; the pixel values in the highlight have been clipped to the maximum pixel value. (b) A dataset image taken on a second pass of the light stage with a 3-stop neutral density filter placed in front of the camera lens. While most of the image is too dark to be useful, the specular highlight in the shell is properly imaged within the range of sensitivity of the camera. The two images can be combined into a single high dynamic range image as in [5]. Relighting the dataset with the method described in Section 4 will produce correct results only if the high dynamic range dataset is used. (c) A rendering of the artifacts using only the low dynamic range imagery captured from a single pass of the light stage as in (a); the abalone necklace piece appears to made of a diffuse plaster-like material. (d) A rendering of the artifacts using high dynamic range imagery captured from multiple passes of the light stage faithfully produces the bright iridescent specular reflection in the necklace.



(a)



(b)



(c)



(d)



(e)

Figure 9: **Capturing Jewelry and Clothing** The light stage can have a chair mounted within it that allows the appearance of traditional clothing and jewelry to be captured as worn by a person, in this case by a member of the particular culture of the artifacts' origin. (a) and (b) show tribal elder George Randall being photometrically recorded in the light stage wearing the otter fur hat, a deerskin Ghostdance shirt, a hair pipe necklace, and holding a prayer staff (Section 8 provides some more detailed information about the artifacts.) Both images used an extended shutter speed of ten seconds to capture an image of all of the lighting directions being illuminated at once. (c) and (d) show close-ups of Randall in the Grace Cathedral and a manually specified lighting environment. (e) shows a wider view of Randall with a prayer staff synthetically re-illuminated in the Grace Cathedral lighting environment.

6 Future Work

A benefit of the technique presented in this paper is that no geometric model of the artifact need be supplied or acquired. However, the technique as so far presented does not allow the viewpoint of the object to be changed, which is the principal way to give a sense of the geometry and reflectance of an object. One technique to render a reflectance field of an artifact from novel viewpoints would be to acquire a geometric model of the artifact, texture map the model with illuminated renderings of the artifact, and then render the texture-mapped model from different points of view. In this paper, we wish to model artifacts whose geometry can not easily be scanned with currently available 3D scanning techniques. As a result, we propose an image-based approach to rendering artifacts from novel viewpoints.

The Light Field [10] and Lumigraph [8] works suggest acquiring a two-dimensional dataset of images from different viewpoints situated on a viewing surface around an object and then composing new viewpoints by sampling similar pixel rays found in existing images. Since this is done entirely with images, no three-dimensional model of the object is required. However, to produce sharp renderings from novel viewpoints the the spacing of the image viewpoints must be very fine, which for our purposes would multiply

the already great quantity of data required to represent an object's reflectance field from just one viewpoint. In order to reduce the amount of data required, a practical suggestion would be to acquire more widely spaced viewpoints and geometrically interpolate between them using view interpolation [20].

To show the potential merit of such an approach, we produced two rendered views (Figures 10(a) and 10(c)) of the headdress made from reflectance fields acquired by two simultaneously running digital video cameras. From these two static viewpoints, we used manually provided image correspondences to produce an intermediate viewpoint as seen in Figure 10(b). Ideally, one would like to find such correspondences manually using an optic flow technique. While finding such correspondences can be challenging [14], we note that the problem may be easier in our case since we have many pictures under different lighting conditions of the object seen from the two viewpoints. Thus, matching could be performed on a much higher-dimensional space of image pixels (under all the different lighting conditionals) to help disambiguate correspondences.

Figure 11 shows a possible augmentation of the light stage apparatus that would capture such datasets. Instead of just a few cameras, a linear array of sixteen cameras aimed at the subject is distributed along a second semicircular arm. Each time the arm




Figure 10: **Rendering the Artifact from Arbitrary Viewpoints (a) and (c)** Renderings of the artifact illuminated by the Grace Cathedral environment from different viewing directions. **(b)** An illuminated rendering of the artifact from an intermediate novel viewpoint made using the image-based rendering technique of view interpolation [20]. Using this technique, the artifact can be rendered from any viewpoint if reflectance field data is captured from a sufficiently dense sampling of discrete viewing directions. An augmented light stage that would automate such a capture process is shown in 11.

of lights goes around, a reflectance field of the artifact is captured from a variety of latitudes but from just one longitudinal direction. To capture views spaced out in longitude as well, there are two possibilities. The first is that the array of cameras would rotate in fixed angular increments after each rotation of the lights; the other is that the artifact itself would be rotated by a motion-controlled platform after each lighting pass.



Figure 11: **An Evolved Light Stage** includes a stationary array of cameras, seen at right, that record the artifact from a one-dimensional array of directions, and a motorized platform that can rotate the artifact a set number of degrees for every rotation of the light stage arm. The device would capture a complete light field of the artifact for every direction of incident illumination, allowing artifacts to be rendered from any viewpoint as well as illuminated from any direction.

In this manner, we could effectively capture an entire light field [10, 8] of the artifact as illuminated from every possible incident illumination direction, and from such a dataset could render the artifact from arbitrary angles and in arbitrary illumination, all without having geometric information for the artifact. For the moment, we leave this acquisition process for future work.

Another important avenue for future work will be to capture higher spectral resolution of the light reflected from the object as illuminated by the light stage. Treating reflectance and incident illumination with just three spectral bands (R, G, and B) is an approximation that can create color rendition problems when either the object's reflectance or the illuminant has a complex spectrum. One

technique for this would be to use a monochrome camera equipped with a multispectral filter wheel to capture an object's reflectance field at a variety of spectral bands. Another technique would be to place a Liquid Crystal Tunable Filter (LCTF) [7] in front of the camera to choose the reflectance bands. Combined with multispectral measurements of incident illumination (taken, for example, by placing an LCTF in front of the camera imaging a mirrored ball), such datasets could yield significantly more accurate renditions of artifacts under novel illumination conditions (such as under fire-light from a torch or fluorescent light in a museum), even if the target rendering space remains just (R,G,B).

7 Conclusion

In this paper we have presented an alternative technique for photometrically acquiring computer graphics models of real-world cultural artifacts. In this work we acquire reflectance fields of the artifacts – image datasets that directly measure how an artifact transforms incident illumination into radiant imagery – rather than surface geometry and texture maps. The method allows the artifacts to be rendered under arbitrary illumination conditions, including image-based illumination sampled from the real world. In this work we have focused on rendering artifacts from the same viewpoints from which their imagery is captured, but have shown that image-based rendering techniques can allow rendering from novel viewpoints to be performed as well. We demonstrated realistic illuminated renderings of a variety of cultural artifacts which would be challenging to model, represent, and render using current digitization techniques. It is our hope that this line of research will eventually help yield practical methods for digitizing photometrically accurate models of cultural artifacts, as well as provide insight into improving current techniques.

8 About the Artifacts

The artifacts featured in this paper are from the private collection of George Randall, a tribal elder from the White Earth Chippewa Reservation in Minnesota.

The headdress in Figure 6 is a classic war bonnet from the Lakota tribe made around 1920 at the Rosemont Reservation in South Dakota near Little Big Horn. The main head band is made from a Union Army blanket, and the headdress features bead work, metal bells, white and black ermine fur, colored ribbons, and turkey feathers augmented with tufts of rabbit fur at the tips. All the materials are organic except for the bells and mirrors, which were typical trade goods.

The drum in Figures 7 and 5 is a ceremonial flat drum made of

stretched elk stomach held together by leather lacing. The painted design shows two eagles, or *hanbli*, a prominent icon in the tribe's narrative tradition.

The otter skin hat in Figure 8 was made in California from a freshwater otter skin from Canada. The hat is approximately twenty years old and is typical of a style of headdress in use for approximately two hundred years.

The choker necklace in Figure 8 is also approximately twenty years old and is made from "hair pipe" (hollowed out bone from small animals), trade beads, and an abalone shell at the front.

The Ghostdance shirt in Figure 9 is sewn from baby deer skin. The front features six strands of horse hair wrapped with red thread, as well as three horse shoe prints indicating that the wearer owns three horses. The back of the shirt features three ermine tails attached to abalone shell circles.

These styles of clothes and jewelry are typical of Native American tribes spanning from the Northeast to Wisconsin such as the Seneca, Mohawk and Wampanoag.

Acknowledgements

We would like to express our gratitude to George Randall for making himself and his collection of artifacts available for this project and our great thanks to Maya Martinez-Smith for organizing his visit and making this collaboration possible. We thank Chris Tchou for his extensive contributions to the software employed in creating the renderings and both Chris Tchou and Dan Maas for writing the interactive reflectance field visualization program. We also wish to thank Brian Emerson for his modeling for figure 11, Andy Gardner for helping with the video editing, and Bobbie Halliday and Jamie Waese for their invaluable production assistance. This work was supported by the University of Southern California, the United States Army, and TOPPAN Printing Co., Inc. and does not necessarily reflect any corresponding positions or policies and no official endorsement should be inferred.

References

- [1] BARIBEAU, R., COURNOYER, L., GODIN, G., AND RIOUX, M. Colour three-dimensional modelling of museum objects. *Imaging the Past, Electronic Imaging and Computer Graphics in Museum and Archaeology* (1996), 199–209.
- [2] CURLESS, B., AND LEVOY, M. A volumetric method for building complex models from range images. In *SIGGRAPH 96* (1996), pp. 303–312.
- [3] DEBEVEC, P. 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).
- [4] DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. Acquiring the reflectance field of a human face. *Proceedings of SIGGRAPH 2000* (July 2000), 145–156. ISBN 1-58113-208-5.
- [5] DEBEVEC, P. E., AND MALIK, J. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97* (August 1997), pp. 369–378.
- [6] DEBEVEC, P. E., TAYLOR, C. J., AND MALIK, J. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *SIGGRAPH 96* (August 1996), pp. 11–20.
- [7] GAT, N. Real-time multi- and hyper-spectral imaging for remote sensing and machine vision: an overview. In *Proc. 1998 ASAE Annual International Mtg.* (Orlando, Florida, July 1998).
- [8] GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. The Lumigraph. In *SIGGRAPH 96* (1996), pp. 43–54.
- [9] HAEBERLI, P. Synthetic lighting for photography. Available at <http://www.sgi.com/grafica/synth/index.html>, January 1992.
- [10] LEVOY, M., AND HANRAHAN, P. Light field rendering. In *SIGGRAPH 96* (1996), pp. 31–42.
- [11] LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., SHADE, J., AND FULK, D. The digital michelangelo project: 3d scanning of large statues. *Proceedings of SIGGRAPH 2000* (July 2000), 131–144. ISBN 1-58113-208-5.
- [12] MALZBENDER, T., GELB, D., AND WOLTERS, H. Polynomial texture maps. *Proceedings of SIGGRAPH 2001* (August 2001), 519–528. ISBN 1-58113-292-1.
- [13] MARSCHNER, S. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, August 1998.
- [14] MCMILLAN, L., AND BISHOP, G. Plenoptic Modeling: An image-based rendering system. In *SIGGRAPH 95* (1995).
- [15] NIMEROFF, J. S., SIMONCELLI, E., AND DORSEY, J. Efficient rendering of naturally illuminated environments. *Fifth Eurographics Workshop on Rendering* (June 1994), 359–373.
- [16] ROCCHINI, C., CIGNONI, P., AND MONTANI, C. Multiple textures stitching and blending on 3d objects. *Eurographics Rendering Workshop 1999* (June 1999). Held in Granada, Spain.
- [17] RUSHMEIER, H., BERNARDINI, F., MITTLEMAN, J., AND TAUBIN, G. Acquiring input for rendering at appropriate levels of detail: Digitizing a pieta. *Eurographics Rendering Workshop 1998* (June 1998), 81–92. ISBN 3-211-83213-0. Held in Vienna, Austria.
- [18] SATO, Y., WHEELER, M. D., AND IKEUCHI, K. Object shape and reflectance modeling from observation. In *SIGGRAPH 97* (1997), pp. 379–387.
- [19] TURK, G., AND LEVOY, M. Zippered polygon meshes from range images. In *SIGGRAPH 94* (1994), pp. 311–318.
- [20] WILLIAMS, L., AND CHEN, E. View interpolation for image synthesis. In *SIGGRAPH 93* (1993).
- [21] WONG, T.-T., HENG, P.-A., OR, S.-H., AND NG, W.-Y. Image-based rendering with controllable illumination. *Eurographics Rendering Workshop 1997* (June 1997), 13–22.
- [22] WOOD, D. N., AZUMA, D. I., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D. H., AND STUETZLE, W. Surface light fields for 3d photography. *Proceedings of SIGGRAPH 2000* (July 2000), 287–296. ISBN 1-58113-208-5.
- [23] Y.CHEN, AND MEDIONI, G. Object modeling from multiple range images. *Image and Vision Computing* 10, 3 (April 1992), 145–155.