Performance Geometry Capture for Spatially Varying Relighting

Andrew Jones Andrew Gardner Mark Bolas† Ian McDowall‡ Paul Debevec

USC Institute for Creative Technologies
University of Southern California†
Fakespace Labs‡

In previous work, we have demonstrated that a traditional lighting basis can be captured for live human performances using a high speed camera and time-multiplexed illumination using LEDs. A linear combination of these bases can be formed to relight the performance under any distant lighting environment. In this paper, we extend this process to simulate locally spatially-varying illumination by including additional structured light patterns from a video projector within the illumination basis. Our processed data is an animated 3D model with both geometric and reflectance information. This augmented dataset can be used to simulate spatially-varying illumination effects. We show how this underlying model can be used to remove indirect illumination from the original lighting basis, and simulate the indirect illumination from the spatially varying light.

Categories and Subject Descriptors: ... [**...**]: ...

General Terms: ...

Additional Key Words and Phrases: ...

---

1. INTRODUCTION

Recent image-based capture systems have allowed researchers and artists to easily capture complex real-world objects and materials. By photographing how an object respond to different lighting conditions or bases, these images can be combined to place the objects under complex novel illumination environments. However, traditional image-based capture and relighting is limited to recreating distant illumination. Full real-world lighting is much broader, including such complex effects as partial shadows and shafts of light.

More complex lighting can be recreated by increasing the number of photographed lighting bases at the expense of overall capture speed. If we wish to capture not only static objects, but moving performances, a constraint is placed on the number of lighting bases that can be included. In this paper we attempt to optimize the set of lighting bases for capturing and recreating complex spatially varying illumination. We propose a hybrid technique for combining directional illumination with additional structured light patterns from a video projector.
The original contributions of our work include:

(1) A system for capturing simultaneous reflectance functions and geometry for real-time performances.

(2) A technique that uses the augmented geometry and reflectance dataset to simulate spatially-varying illumination effects in an image-based relighting process.

(3) A technique for modifying the image-based renders to compensate indirect effects that occur under spatially-varying illumination.

2. PREVIOUS WORK

2.1 Reflectance Capture

[Debevec et al. 2000] introduced the *reflectance field* concept as an 8D reflectance function. The reflectance field relates the outgoing reflectance (u_r, v_r) in any given direction to the incoming flux (u_i, v_i) in another given direction in the following form:

$$R = R(R_i, R_r) = R(u_i, v_i, \theta_i, \phi_i, u_r, v_r, \theta_r, \phi_r) \tag{1}$$

[Debevec et al. 2000] sampled the reflectance field with a single moving directional light source and a camera placed at several locations. Due to time and storage limitations, the original light stage device did not position the light and camera along every possible direction. As a result the system only captured 4D non-local reflectance fields with a limited viewpoint and distant incident illumination. More recent work [Hawkins et al. 2001] [Hawkins et al. 2004] improved capture speed by increasing the number of moving lights. [Wenger et al. 2005] were able to capture a full 4D data set at real-time rates using a high speed camera and a static geodesic dome of time-multiplexed LEDS. Alternatively [Matusik et al. 2002] incorporated a moving camera array into a light stage to capture a 6D reflectance function with additional viewing dimensions.

All these image-based relighting techniques are *not* capable of simulating spatially-varying incident illumination, such as dappled illumination or partial shadows on the subject, since no basis illumination condition shows such effects.

Local 6D reflectance fields have previously been captured by expanding the lighting basis to include spatially-varying conditions. [Masselus et al. 2003] replaced the moving point light source with a projector but at the expense of greatly increased capture times. [Zen et al. 2005] took advantage of Helmholtz reciprocity to measure multiple incident lighting directions simultaneously. While this provided a significant speed increase over [Masselus et al. 2003], the adaptive processing and large number of projected patterns still required capture times to be on the order of 5 minutes per frame.

In this report, we present a hybrid process for simulating spatially-varying illumination on 4D datasets such as those of [Debevec et al. 2000] [Hawkins et al. 2001] [Wenger et al. 2005] by including additional structured light patterns from a video projector within the illumination basis. This technique allows us to achieve 6D lighting flexibility at speeds associated with previous 4D capture setups.

2.2 Geometry Capture

While previous work has not been able to capture real-time animated 6D reflectance fields, real-time animated geometry has been achieved by several researchers.

[Rusinkiewicz et al. 2002] extended structured light coding schemes to reconstruct moving rigid objects. Similarly [Sa et al. 2002] used coded stripe boundaries with complimentary colors to recover moving geometry and a simultaneous video texture but only illuminated from the projector position. [Li Zhang 2004] captured real-time faces using stereo and rapidly projected video patterns. None of these works however focused on reproducing the reflectance properties of the subject or simulating spatially-varying illumination.

[Debevec et al. 2000], captured face geometry using binary structured light patterns with a corresponding reflectance field. However the speed of the capture device limited the subject to static poses. In this work we use similar patterns to capture geometry but displayed using a high-speed projector.

3. DATA CAPTURE

Our geometry and reflectance datasets are captured 24 times per second using a Luxeon V LED-based light stage, a Vision Research high-speed Phantom 7.1 camera, and a specially modified 1024×768 pixel high-speed DLP projector. Details on the lighting apparatus can be found in [Wenger et al. 2005]. The lights, camera, and projector are synchronized at approximately 1500 frames per second. For each 24th of a second, the subject is lit by a series of 29 basis lighting directions (e.g. Fig 1) followed by 24 structured light patterns (e.g. Fig 2).

We use simple binary structured light patterns [Tchou 2002] for ease of implementation. At the high frame-rate used by our camera, subject motion is minimal. A more advanced coding scheme that compensates for motion as as that seen in [Rusinkiewicz et al. 2002] or [Sa et al. 2002]could easily be used.

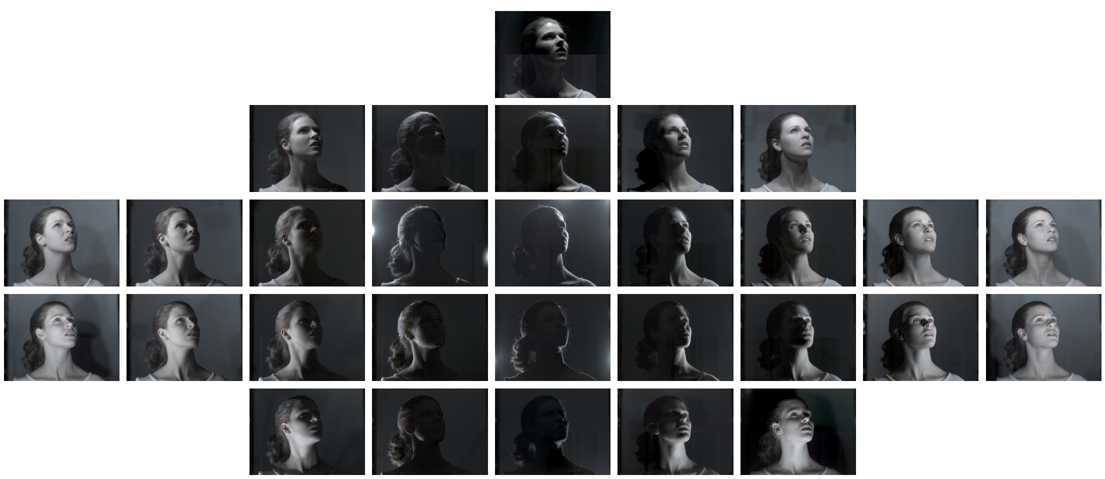

Fig. 1. 29 basis lighting directions for a given frame

A surface mesh for each frame of the performance is recovered from the structured lighting patterns. As we capture surface reflectance function, we can also recover surface normals and spatially-varying diffuse color (albedo) using the photometric stereo based techniques described in [Wenger et al. 2005]. The results are shown in Figure 3. We use a gradient-descent optimization to refine the geometry to match the estimated surface normals. Alternative mesh optimization methods such as [Nehab et al. 2005] could also be used.

Fig. 2. 24 binary structured light patterns for a given frame

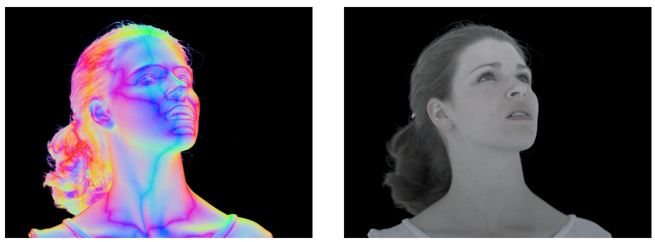

(a) (b)

Fig. 3. Photometric stereo (a) Color-coded normals (b) Albedo

4. RELIGHTING

In image-based relighting, novel illumination on a subject is synthesized based on images acquired in different basis lighting conditions. Most commonly, the basis images of the subject are taken under a variety of directional lighting conditions. A linear combination of the color channels of the photographed lighting directions is formed to produce the rendering under novel illumination. When the lighting directions are densely distributed throughout the sphere, any distant lighting environment can be simulated accurately. We apply this technique to our lighting bases to simulate distant illumination (Fig. 4). The limitation of pure image-based rendering to the lighting and viewing dimensions spanned by the basis images.

An alternative approach is driven by the 3d geometric model. Most commonly a parametric reflectance function or BRDF is fit to the captured image data. This technique has been used in [Debevec et al. 2000] and [Lensch et al. 2001] to arbitrarily change the camera viewpoint and lighting enviromnent. The disadvantage of this approach is the recovered BRDF often fails to capture subtle or complex reflectance effects. Robust reflectance fitting also becomes difficult when the viewing and lighting directions are sparsely sampled.

Our approach is closer to that of [Debevec 1998] in that it works through differential rendering. This method creates a 3d geometric and reflectance model, but only uses it guide modifications to the initial image-based rendering.

For traditional image-based rendering, the primary assumption is that incident illumination is spatially consistent throughout the capture volume. This assumption is what allows a constant scale to be applied to each entire basis image. This assumption holds for distant illumination as the subject is small in relation to the

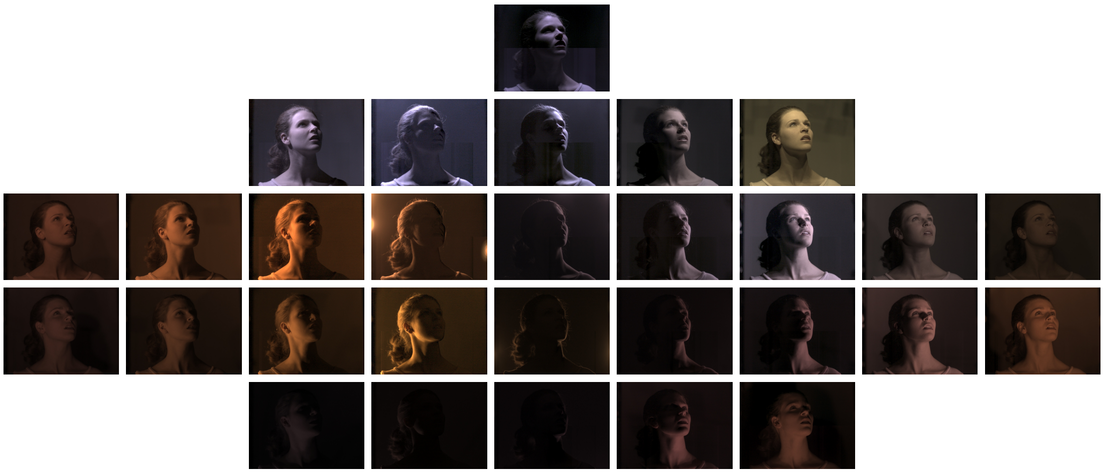

Fig. 4. Basis lighting directions scaled by sampled light probe in Grace Cathedral

lighting environment. If a light source is placed close to the subject or partially obscured so the subject is in partial shadow, this assumption is no longer valid. A *spatially-varying light*is one where the lighting varies throughout the capture volume and across the rendered image. For such a light source, each point in the image must be scaled independently.

The geometric model allows us to determine where each point in the image lies within the three-dimensional volume of a spatially-varying light. As the subject moves through the varying light field, we can find the appropriate illumination that correctly follow the contours of the face. We simulate this effect using a photograph of an HDR stained glass window (Fig. 5) to modulate one of basis lighting directions (Fig. 6 (a)). In this case, the lighting lookup become a simple projective texture map (Fig. 6 (b)). When the initial basis image is multiplied by the projected texture, the result is an approximation of spatially varying illumination (Fig. 6 (c)).

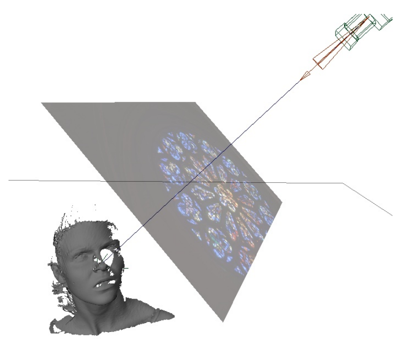

Fig. 5. An HDR image of a stained glass window generates a spatially-varying point light source.

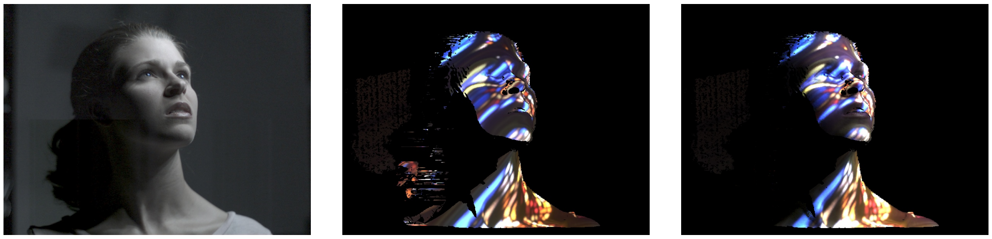

(a) (b) (c)

Fig. 6. (a) Original directional lighting basis, (b) Projected spatially-varying light, (c) Naive relighting without compensation for pixel-interdependencies (a x b)

However, this initial process does not correctly model second-order illumination effects. For example, if a spatially-varying light source illuminates just a person's shoulder, in reality this would produce indirect illumination on the underside of the chin. In Figure 6 (c), this is most notable as missing bounce light under the chin and incorrectly scaled bounced light around the eyes. Correcting for this effect requires factoring indirect illumination from direct illumination in the initial lighting basis. As we do not independently observe the indirect illumination, we instead simulate it using our 3D model texture-mapped by the recovered albedo (Fig. 3 (b)). This indirect simulation is then used to partially correct the naive technique shown in Figure 6.

For a given basis lighting direction, we match the illumination produced by the light stage, and use it to light the recovered geometry. This allows us to estimate and subtract the existing bounced light from the basis image (Fig. 7(c)). An indirect illumination pass can be computed in most standard global illumination render engines either by directly caching bounced irradiance, or by comparing two renders of the diffuse geometry, one with a simulated indirect bounce and one without. To simulate direct lighting effects, we multiply the resulting image by same projected spatially varying illumination used above (Fig. 6(b)) to get (Fig. 7(d)). Finally, we add new simulated indirect illumination (Fig. 7(f)) generated by the projected illumination. The final rendering combines spatially-varying illumination with fill light from a conventional light probe (Fig. 8).

5. CONCLUSION

We have demonstrated our technique on a four-second sequence of an actor's face shown in the video. Selected frames are shown in Figure 8. To our knowledge this is the first sequence capturing time-varying geometry and directionally-varying surface reflectance, and the first simulation of indirect illumination effects from a traditional image-based relighting dataset.

This differential rendering technique has the benefit of preserving complex surface reflection properties for direct surface reflection, while using the diffuse estimate solely for indirect effects. While traditional photometric stereo estimates only diffuse reflectance component, is possible to fit more complicated BRDF to the reflectance functions as in [Lensch et al. 2001] for more accurate indirect reflections.

In our current work, we are extending this process to multiple projectors to capture more complete subject geometry and to more accurately simulate subsurface

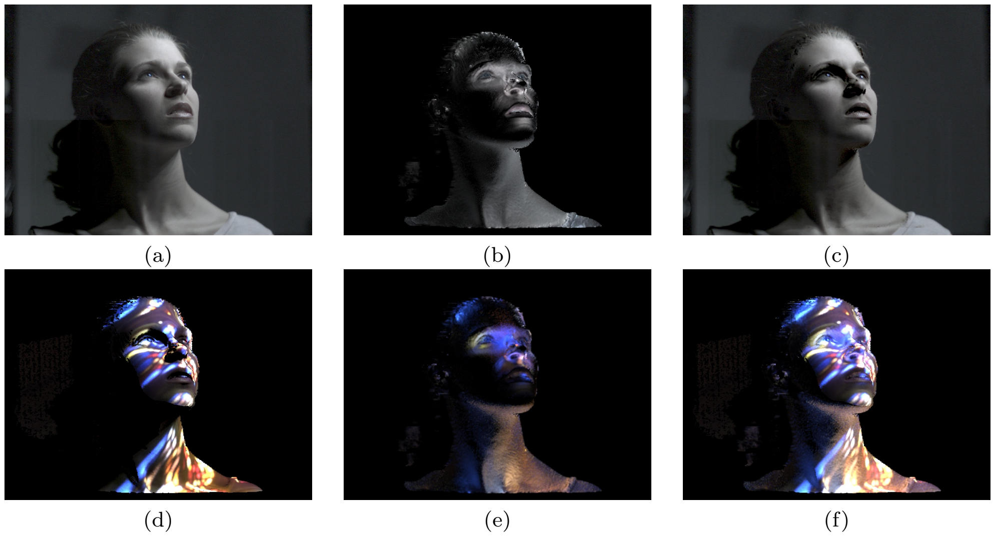

Fig. 7. (a) Original directional lighting basis, (b) Simulated bounced light component of (a), (c) Basis image with bounced light removed (a - b), (d) Image c scaled by projected spatially-varying light (e) Simulated bounced light for d, (f) Rendering with spatially varying light including both direct and indirect illumination effects (d + e)

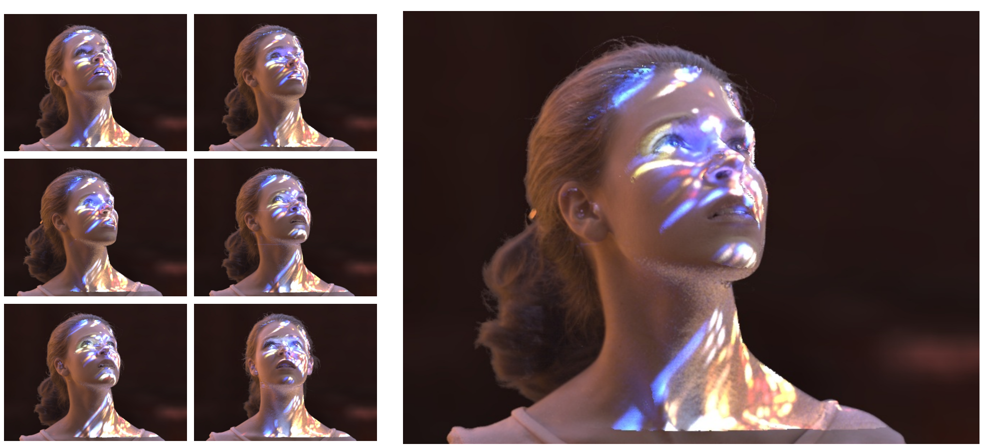

Fig. 8. The performance relit by Grace Cathedral in San Franscisco. The final composite includes both local and distant illumination.

scattering effects resulting from spatially-varying illumination.

REFERENCES

DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of SIGGRAPH 98*.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000.

Acquiring the reflectance field of a human face. In *Proceedings of ACM SIGGRAPH 2000*. Computer Graphics Proceedings, Annual Conference Series. 145–156.

Hawkins, T., Cohen, J., and Debevec, P. 2001. A photometric approach to digitizing cultural artifacts. In *Proc. 2nd International Symposium on Virtual Reality, Archaeology, and Cultural Heritage (VAST 2001)*. 333–342.

Hawkins, T., Wenger, A., Tchou, C., Gardner, A., Göransson, F., and Debevec, P. 2004. Animatable facial reflectance fields. In *Rendering Techniques 2004: 15th Eurographics Workshop on Rendering*. 309–320.

Lensch, H. P. A., Kautz, J., Goesele, M., Heidrich, W., and Seidel, H.-P. 2001. Image-based reconstruction of spatially varying materials. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*. 103–114.

Li Zhang, Noah Snavely, B. C. S. S. 2004. Spacetime faces: High-resolution capture for modeling and animation. *ACM Transactions on Graphics (SIGGRAPH 2004)*.

Masselus, V., Peers, P., Dutré, P., and Willems, Y. D. 2003. Relighting with 4d incident light fields. *ACM Transactions on Graphics (SIGGRAPH 2003) 22,* 3 (July), 613–620.

Matusik, W., Pfister, H., Ngan, A., Beardsley, P., Ziegler, R., and McMillan, L. 2002. Image-based 3d photography using opacity hulls. *ACM Transactions on Graphics 21,* 3 (July), 427–437.

Nehab, D., Rusinkiewicz, S., Davis, J., and Ramamoorthi, R. 2005. Efficiently combining positions and normals for precise 3d geometry. *ACM Transactions on Graphics (SIGGRAPH 2005)*.

Rusinkiewicz, S., Hall-Holt, O., and Levoy, M. 2002. Real-time 3d model acquisition. *ACM Transactions on Graphics 21,* 3 (July), 438–446.

Sa, A., Carvalho, P. C. P., and Velho, L. 2002. (b,s)-bcsl : Structured light color boundary coding for 3d photography. In *Proceedings of 7th International Fall Workshop on VISION, MODELING, AND VISUALIZATION*.

Tchou, C. D. X. N. 2002. Image-based models: Geometry and reflectance acquisition systems. M.S. thesis, University of California at Berkeley.

Wenger, A., Gardner, A., Unger, J., Tchou, C., Hawkins, T., and Debevec, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics (SIGGRAPH 2005)*.

Zen, P., Chen, B., Garg, G., Marschner, S., Horowitz, M., Levoy, M., and Lensch, H. 2005. Dual photography. *ACM Transactions on Graphics (SIGGRAPH 2005)*.

...