

Performance Geometry Capture for Spatially Varying Relighting

Andrew Jones Andrew Gardner
USC Institute for Creative Technologies

Mark Bolas† Ian McDowall‡
University of Southern California†

Paul Debevec
Fakespace Labs‡

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 different directional lighting conditions. A linear combination of the color channels of the different 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.




Figure 1: (a) One of 29 basis lighting directions (b) One of 24 structured light patterns (c) Indirect illumination computed from the recovered geometry and albedo (d) Rendering with spatially varying light including both direct and indirect illumination effects.

This standard image-based relighting technique is *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. [Masselus et al. 2003] expanded the lighting basis to include spatially-varying conditions, but at the expense of greatly increased capture times. Recent work [Wenger et al. 2005] has shown that a traditional lighting basis can be captured for live human performances using a high speed camera and time-multiplexed illumination using LEDs. In this sketch, we present a process for simulating spatially-varying illumination on such datasets 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. We show how this augmented dataset can be used to simulate spatially-varying illumination effects in an image-based relighting process.

Our geometry and reflectance datasets are captured 24 times per second using a Luxeon V LED-based light stage, a Vision Research high-speed camera, and a specially modified 1024×768 pixel high-speed DLP projector. The lights, camera, and projector are synchronized at 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(a)) followed by 24 structured light patterns (e.g. Fig. 1(b)).

We recover surface normals and spatially-varying diffuse color (albedo) using the photometric stereo based techniques described in [Wenger et al. 2005]. A surface mesh of the geometry is recovered from the structured lighting patterns. We use a gradient-descent optimization to refine the geometry to match the estimated surface normals. The normals are also used to partially fill in missing geometry resulting from projector occlusions.

The geometric model allows us to determine where each point in the image lies within the three-dimensional volume of a spatially-varying light source, such as a projected stained glass window (Fig. 2). With this information, each point on the face can be scaled as if illuminated by a different lighting environment, giving the appearance of spatially-varying illumination. However, this 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. We have developed a process for simulating such effects by performing a global illumination computation on the 3D model texture-mapped by the recovered albedo. For a given basis lighting direction, we estimate and remove the existing bounced light in the image. We multiply the resulting image by the projected spatially varying illumination. Finally, we add new simulated indirect illumination (Fig. 1(c)). The final rendering combines spatially-varying illumination with fill light from a conventional light probe (Fig. 1(d)).

We have demonstrated our technique on a four-second sequence of an actor's face, shown in the video. 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. In our current work, we are extending this process to multiple projectors to capture more complete subject geometry and to more accurately simulate subsurface scattering effects resulting from spatially-varying illumination.




Figure 2: An HDR image of a stained glass window is used to create a spatially-varying point light source.

References

- MASSELUS, V., PEERS, P., DUTRÉ, P., AND WILLEMS, Y. D. 2003. Relighting with 4d incident light fields. *ACM Transactions on Graphics* 22, 3 (July), 613–620.
- WENGER, A., GARDNER, A., UNGER, J., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *Conditionally Accepted to ACM Transactions on Graphics (SIGGRAPH 2005)*.