Direct HDR Capture of the Sun and Sky

Jessi Stumpfel Andrew Jones Andreas Wenger
Chris Tchou Tim Hawkins Paul Debevec *

University of Southern California Institute for Creative Technologies†

ABSTRACT

We present a technique for capturing the extreme dynamic range of natural illumination environments that include the sun and sky, which has presented a challenge for traditional high dynamic range photography processes. We find that through careful selection of exposure times, aperture, and neutral density filters that this full range can be covered in seven exposures with a standard digital camera. We discuss the particular calibration issues such as lens vignetting, infrared sensitivity, and spectral transmission of neutral density filters which must be addressed. We present an adaptive exposure range adjustment technique for minimizing the number of exposures necessary. We demonstrate our results by showing time-lapse renderings of a complex scene illuminated by high-resolution, high dynamic range natural illumination environments.

Categories and Subject Descriptors

I.4.1 [**Image Processing and Computer Vision**]: Digitization and Image Capture; I.3.7 [**Computing Methodologies**]: Computer Graphics Three-Dimensional Graphics and Realism

General Terms

Measurement, Lighting, Photometry

Keywords

Image-based lighting, Real-world capture, High dynamic range photography

1. INTRODUCTION

A significant area for recent research has been in digitizing the geometry, reflectance properties, and illumination of real-world scenes to create more realistic computer renderings. Recent techniques employing omnidirectional high dynamic range photography have produced useful datasets of real-world illumination environments [5], which have been used as sources of illumination for both offline [4] and real-time [18, 19] rendering techniques. While datasets seen so far include both indoor and outdoor lighting environments, current techniques have not been able to record an outdoor environment that includes a directly visible sun, which is a common and important type of environment to capture.

Figure 1: The camera with 180° fisheye lens and laptop on the roof of our building (left). This is a typical assembled probe (right) taken with the fisheye lens and neutral density filter spanning the 17 stops of the sky and sun.

Capturing the upper hemisphere of the sky is easily done using a fisheye lens, but there are two principal challenges in capturing the full dynamic range of outdoor illumination. The first is the breadth of the range - the sun can be well over five orders of magnitude (or seventeen stops) brighter than the sky and clouds, which is a greater range than can be covered with typically available shutter speeds. The second is the absolute intensity of the sun, which is much brighter than what cameras are designed to capture. In this paper, we overcome these problems by using a combination of varying shutter speed, aperture, and a properly chosen neutral density filter, and we describe the calibration procedures necessary to make use of all three of these light attenuation mechanisms. With optimized settings, we are able to capture the full dynamic range of the sky using at most seven exposures using a camera with a 12-bit linear-response sensor. In this work we also discuss an adaptive capture technique in which fewer exposures are acquired when the sun is not visible.

To demonstrate these techniques we captured several full days of natural outdoor illumination which we use to render an outdoor scene under time-lapse illumination conditions.

*{stumpfel,jones,wenger,tchou,timh,debevec}@ict.usc.edu
†13274 Fiji Way, Marina Del Rey, CA 90292

.

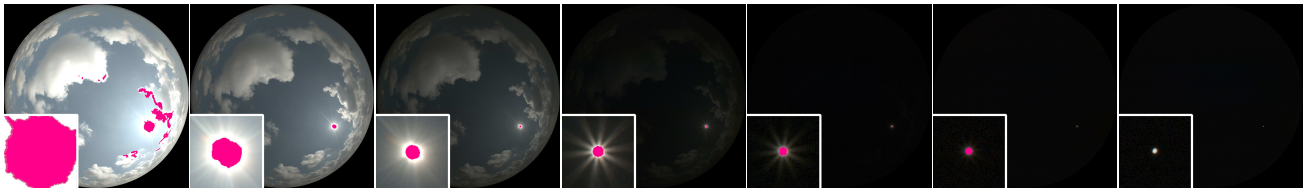

Figure 2: HDR sequence and camera settings to span the 17 stops of the sky and sun in 7 exposures. A detailed view of the sun is shown in the bottom left of each image; pink regions indicate saturated pixels. The darkest image is the only image that does not saturate for the sun.

2. RELATED WORK

To date, most computer graphics renderings illuminated by sun and sky light have been done using analytic sky models as in [16, 21] and the "gensky" function in [22]. A model specifically designed for realistic rendering of large-scale environments was presented in [17]; this model used an atmospheric scattering model to compute not only the colors in the sky but also physically-based aerial perspective of surfaces near the ground. A night sky model was developed by [11] where they model the effect of the sun, moon, atmosphere, orbiting dust, stars, and the Milky Way. [23] adapted an analytic sky model so that it could be fit to a small set of sky intensity measurements for an inverse rendering application. Work has also been done to efficiently compute the lighting of outdoor natural illumination on synthetic scenes such as [3]. These models can produce synthetic skies that match ideal conditions well, but they do not model the more visually interesting aspects of cloud formation and motion. Several techniques have been proposed to simulate cloud formation, illumination, and evolution [9, 7, 15]. These techniques produce synthetic 3D clouds that are realistic in appearance, but do not necessarily represent the full range of atmospheric effects encountered in nature, or the specific types of cloud formation found in particular geographic locations.

In the meteorological community, instruments and cameras monitor weather year round at locations worldwide. The goal in this case is to compile weather statistics, such as average cloud cover and precipitation; one such study is presented by [14]. These statistics do not require capturing the full dynamic range of the sun and sky. Instead they use a moving arm to block the view of the sun from the sensor.

In our work, we use an image-based technique to directly image the sky, clouds, and sun based on the illumination capture process of [4], adapted to capture the full dynamic range of the sky up to and including the sun. A technique proposed in [6] reconstructs the intensity of the sun by observing a diffuse sphere in addition to a mirrored sphere used to capture the sky; however, this fails to capture brightness detail in the circumsolar region and is less useful for rendering direct views of the sky. Finally, we use our technique to capture time-lapse natural illumination measurements throughout the day in a manner similar to the high dynamic range video project of [12].

3. CAPTURING THE FULL HDR RANGE OF THE SUN AND SKY

The roof of our laboratory provided an unobstructed view of the horizon in all directions. The entire sky could be imaged from this vantage point by placing a Canon EOS 1DS camera pointing up, equipped with a 8mm Sigma fisheye lens [Figure 1]. While the Canon camera provides automatic exposure bracketing, using this to adjust the exposure between 2 seconds and 1/8000th of a second yields only 14 stops. Shutter speed times longer than 2 seconds are impractical because of the motion of the clouds. Since each stop changes the light level by a factor of 2, that is a dynamic range of 2^{14} (or 16384). This is insufficient to capture the dynamic range of a typical sunny sky.

A library provided by Canon [1] allowed us to control the camera from a laptop, while downloading the images to the laptop's disk. This allowed us to programmatically adjust the aperture and exposure, acquiring an image sequence detailed in Figure 2. This sequence spanned the needed 17 stops in a manner sufficient for reconstructing the HDR image, and required approximately 50 seconds to acquire and download each series. We consider each individual image to have a dynamic range of 5 stops. This 5 stop range in addition to the 17 stop variance in the sequence, provides a 22 stop range, or over 4 million times the light. Our images are taken at most 3 stops apart, allowing sufficient dynamic range overlap between images to assemble them into an HDR image.

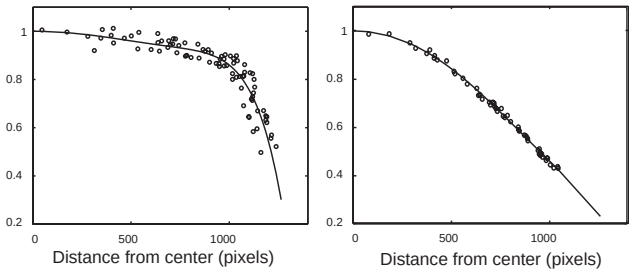

Figure 3: Graphs representing relative pixel brightness for given distances from the image center. Circles represent measurements of a reference illuminant at image locations. Curved lines indicate the fitted polynomial function. Falloff curve for Sigma 8mm lens at f/16 (left) and falloff curve for the same lens at f/4 (right).

3.1 Calibration

To achieve accurate results, it was necessary to calibrate the fisheye lens geometrically and photometrically. Geometric fisheye calibration was accomplished by marking scene

correspondences while rotating the camera around the average nodal point. We assume an ideal fisheye projection [10], and solve for image center and radius of the lens. A more sophisticated geometric model could be based on a more accurate measurement of the lens construction [13], or by fitting a lens distortion model. Also, there is significant radial intensity falloff and vignetting associated with fisheye lens. This has two causes: first, exposure falloff across the image sensor increases with wider apertures, and secondly at larger angles occlusion due to other lens elements introduces vignetting [13]. As seen in Figure 3, we fit a separate radial polynomial function to the data for both f/4 and f/16 apertures [6].

Given the sensitivity of our camera's image sensor, we found that imaging the sun without saturation required at least a 3.0 Neutral Density (ND) filter, allowing just 0.1% of the light to pass through. ND filters are usually placed in front of the lens, but no standard filters fit the curved front surface of the Sigma 8mm fisheye lens. Instead we placed a piece of a Kodak WRATTEN 3.0 Neutral Density gelatin filter between the lens and the camera.

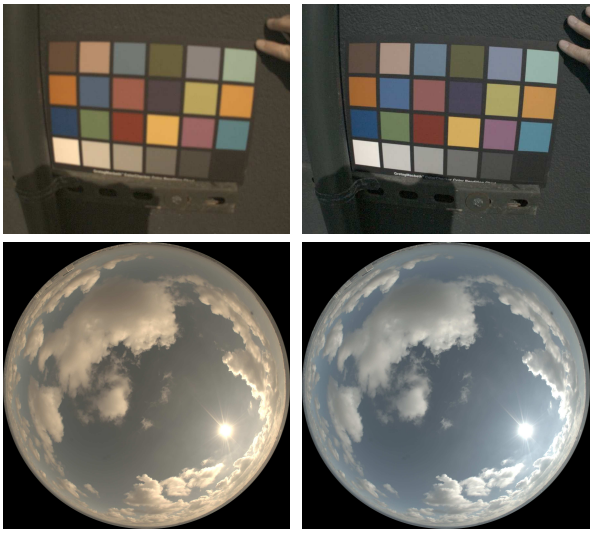

Figure 4: The neutral density filter while reducing the amount of light introduces a chromatic change as well. To correct for this a color chart was photographed with and without the ND filter. These images were then used to solve for a color correction matrix.

Figure 4 shows an uncorrected image of the sky with reddish-brown clouds taken with the ND filter. Contrary to their name, neutral density filters are typically not chromatically neutral. We estimated a correction for the chromatic shift in the visible portion of the ND filter spectrum. We photographed a Macbeth color chart in sunlight with and without the ND filter [Figure 4], then solved for a 3x3 linear color transform that most closely matched the color samples between the images. This transforms the images into the raw color space of the Canon EOS 1DS camera, which could then be mapped into a standardized color space.

To further investigate the effect of the ND filters, we used a Photo Research PR-650 spectroradiometer to measure the ND filter's spectral transmissivity [Figure 5]. Not only is the

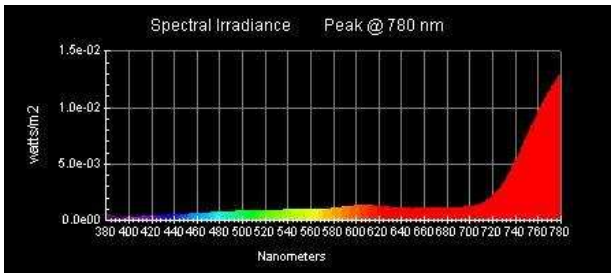

Figure 5: The spectral calibration of the 3.0 ND Kodak WRATTEN Filter, showing a gentle increase in the visible range from blue to red. A significant amount of IR light is transmitted.

filter more transmissive in the red region than the blue, but it also displays a large amount of light transmitted in the near infrared region above 720nm, which not all ND filters are designed to block.

This can be potentially problematic, as many digital cameras have a significant sensitivity to near IR light. Attenuating the visible light can only increase the error resulting from this sensitivity. To determine the influence of IR light we took multiple photographs under an incandescent illuminant. By taking measurements with and without an IR-pass filter, that removes visible wavelengths, we determined the contribution of IR light to the recorded pixel values [Table 1]. Fortunately, our experiments showed that our Canon E0S 1DS does not have a significant near IR response. However other cameras (eg. the Sony VX1000) exhibit strong IR response. As sunlight has a very strong near IR component, any camera that exhibits IR response should use an IR-cutoff filter.

IR Contribution - Sony VX1000

Filter	R	G	B
none	0.74%	0.03%	0.08%
3.0 ND	6.02%	1.00%	3.25%

IR Contribution - Canon D0S 1DS

Filter	R	G	B
none	0.007%	0.006%	0.010%
3.0 ND	0.23%	0.29%	0.57%

Table 1: Percentage of sensor response caused by incandescent infrared light for the Canon EOS 1DS and the Sony VX1000.

4. ADAPTIVE HDR IMAGING

The f/4 1 second exposure was not long enough to capture dark clouds before and during dawn and after dusk. In general, over the course of the day the sky brightness can change drastically with sun position and weather.

To address this problem, as well as to increase the speed of our capture and reduce wear on the camera shutter, we created a process to adaptively select the exposures and apertures when capturing an HDR sequence. This program analyzes each image as it is downloaded to determine when pixels in the sky are underexposed or saturated. The following pseudo-code describes this adaptive capture technique.

```
SetLightMetering(CenterWt_Avg);
while (captureHDRSequences) {
   imageA = TakePhoto(F4, Av);

   tooDark = imageA.IsUnderExposed();
   tooBright = imageA.IsSaturated();

   exposeTime = imageA.GetTv() * 8;
   while (tooDark) {
        image = TakePhoto(F4, exposeTime);
        exposeTime = exposeTime * 8;
        tooDark = image.IsUnderExposed();
    }

    exposeTime = imageA.GetTv() * 2;
    while (tooBright) {
        image = TakePhoto(F16, exposeTime);
        exposeTime = exposeTime / 8;
        tooBright = image.IsSaturated();
    }
}
```

The program captures an HDR sequence by first capturing an image in aperture priority mode (Av). The light metering on the camera is set to "Center-weighted Average Metering". This allows the camera to choose a shutter speed that properly exposes the majority of the image. This image is then analyzed to determine if shorter and/or longer exposures are required. Taking into account both aperture and exposure time, the images in our adaptive HDR sequences are always three stops apart.

We classify an image as underexposed and/or saturated by comparing the raw pixel values against an upper and lower threshold. This method was makes the analysis as fast as possible in order to maximize capture speed. The raw pixel values on Canon EOS class cameras provide 12 bits of data, exhibiting a range from 0 to 16380 in steps of 4. We have measured the Canon sensors to be linearly responsive to light up to a pixel value of 12000. At this point the pixels become less sensitive, assumably due to saturation. To avoid problems of non-linearity, we set our saturation threshold at 9000 and our underexposure threshold at 500.

Using the adaptive capture program we were able to capture an HDR image every 40 seconds with sequences consisting of three to seven pictures. The adaptive capture also allowed us to handle the highly increasing and decreasing light levels during the entire day from pre-dawn to post-dusk. The primary speed bottleneck we encountered was file transfer from the camera to the laptop, which will hopefully improve in new camera systems.

Motion artifacts appear primarily where bright fast moving clouds span multiple exposures. Away from the sun, clouds are easily captured in fewer than three exposures. Chromatic abberations from cloud movement can be reduced during HDR assembly by color-interpolating each photograph in advance.

5. RENDERING AND RESULTS

The captured lighting presented here could be used to render any outdoor scene providing realistic illumination as well as a sky backdrop. Due to the extreme contrast between the sun and sky, the standard global illumination algorithms will produce a large amount of noise, as they will have difficulty sampling the sun sufficiently.

Algorithms have been developed for intelligent sampling of the environment with priority given to high intensity regions such as in [2]. We utilize a simpler solution, approximating the sun as a 0.53 degree diameter directional light, while using traditional sampling techniques for the remainder of the sky. Pixels in the lighting environment that are above a high threshold are set to black and represented by a directional light; this directional light is placed at the weighted centroid of the blackened pixels and set to their total energy. By manually adjusting the threshold, we select a balance between accuracy and noise.

With this method we rendered a synthetic scene, illuminated by an entire day of captured light, using the Arnold global illumination system [8]. Several frames are shown in Figure 7. For these frames, the threshold was set to the peak cloud intensity in the sequence. These datasets have allowed us to visualize the full dynamic range of natural illumination over the course of the day, as seen in Figure 6.

5.1 Future Work

In the future, we would like to extend optical flow algorithms to interpolate between extreme-high dynamic range frames and create smoother time-lapse animations [12]. Motion compensation is difficult as the clouds are not visible in several photographs. We would also like to look into possibilities for removing lens flare artifacts observed in the fisheye images. Because skylight is partially polarized we are exploring techniques for capturing this polarization [20]. Methods of spectral recovery in natural illumination environments also merit investigations.

6. CONCLUSION

Using these techniques we have captured several days of light with variable weather conditions. We believe these HDR images are the first calibrated recordings of natural illumination in high- resolution including a visible sun. Exploring the practical issues of capturing the full high dynamic range of the sky using a fisheye lens with a neutral density filter allows the natural dynamics of the sky and sun to be represented faithfully in our renderings. Several captured HDR days of light can provide flexibility for choosing lighting for novel outdoor renderings. Datasets taken for this paper are available at: http://www.ict.usc.edu/graphics/skyprobes/.

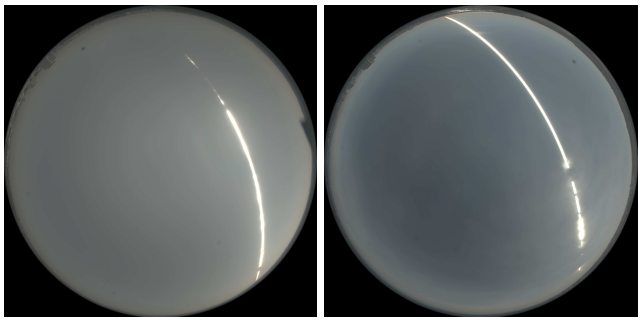

Figure 6: These images are formed as the average of 700 HDR images of the sky taken on February 19, 2004 (left) and February 23, 2004 (right) at one minute intervals. The sun streak is occluded at times by clouds over the day, and individual clouds are averaged to a constant color.

Figure 7: Rendering of a virtual model of the Parthenon with lighting from 7:04am (top left), 10:35am (top right), 4:11pm (bottom left), and 5:37pm (bottom right). Capturing high-resolution outdoor lighting environments with over 17 stops of dynamic range with time lapse photography allows for realistic lighting and effects such as simulated lens post-processed sun flare, as seen in the top left.

Acknowledgments

We would like to thank Lora Chen, David Wertheimer, Richard Lindheim, and Neil Sullivan. Also special thanks to Richard DiNinni for disarming the roof alarm when necessary. This work has been sponsored by the University of Southern California and U.S. Army contract number DAAD19-99-D-0046. Any opinions, findings, and conclusions or recommendations expressed in this paper do not necessarily reflect the views of the sponsors.

7. REFERENCES

[1] Canon digital camera software developers kit, 2004. http://consumer.usa.canon.com/.

[2] S. Agarwal, R. Ramamoorthi, S. Belongie, and H. W. Jensen. Structured importance sampling of environment maps.

[3] K. Daubert, H. Schirmacher, F. X. Sillion, and G. Drettakis. Hierarchical lighting simulation for outdoor scenes. In *Eurographics Rendering Workshop 1997*.

[4] P. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. SIGGRAPH'98*.

[5] P. Debevec. Light probe image gallery, 1999. http://www.debevec.org/Probes/.

[6] P. Debevec et al. Estimating surface reflectance properties of a complex scene under captured natural illumination. *Conditionally Accepted to ACM Transactions on Graphics*, 2004.

[7] Y. Dobashi, K. Kaneda, H. Yamashita, T. Okita, and T. Nishita. A simple, efficient method for realistic animation of clouds. In *Proc. SIGGRAPH 2000*.

[8] M. Fajardo. Monte carlo ray tracing in action. *ACM SIGGRAPH 2001 Course 29*.

[9] G. Y. Gardner. Visual simulation of clouds. In *Proc. SIGGRAPH'85*.

[10] N. Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics and Applications*, 1986.

[11] H. W. Jensen, F. Durand, M. M. Stark, S. Premoze, J. Dorsey, and P. Shirley. A physically-based night sky model. In *Proc. SIGGRAPH 2001*.

[12] S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High dynamic range video. In *Proc. SIGGRAPH 2003*.

[13] C. Kolb, D. Mitchell, and P. Hanrahan. A realistic camera model for computer graphics. In *Proc. SIGGRAPH'95*.

[14] C. Long and J. Deluishi. Development of an automated hemispheric sky imager for cloud fraction retrievals. In *10th Symposium on Meteorological Observations and Instrumentation*, Boston, MA, 1998. American Meteorological Society.

[15] R. Miyazaki, S. Yoshida, Y. Dobashi, and T. Nishita. A method for modeling clouds based on atmospheric fluid dynamics. In *9th Pacific Conference on Computer Graphics and Applications*, 2001.

[16] T. Nishita and E. Nakamae. Continuous tone representation of three-dimensional objects illuminated by sky light. In *Proc. SIGGRAPH'86*.

[17] A. J. Preetham, P. S. Shirley, and B. E. Smits. A practical analytic model for daylight. In *Proc. SIGGRAPH'99*.

[18] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. In *Proc. SIGGRAPH 2002*.

[19] P. Sloan, J. Kautz, and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *Proc. SIGGRAPH 2002*.

[20] J. Stumpfel. Outdoor lighting capture for rendering and inverse global illumination. Master's thesis, California Institute of Technology, 2004.

[21] K. Tadamura, E. Nakamae, K. Kaneda, M. Baba, H. Yamashita, and T. Nishita. Modeling of skylight and rendering of outdoor scenes. *Comput. Graph. Forum*, 12(3), 1993.

[22] G. J. Ward. The RADIANCE lighting simulation and rendering system. In *SIGGRAPH'94*.

[23] Y. Yu and J. Malik. Recovering photometric properties of architectural scenes from photographs. In *Proc. SIGGRAPH'98*.