

Virtual Cinematography: Relighting through Computation

Paul Debevec

USC Centers for Creative Technologies

Recording how scenes transform incident illumination into radiant light is an active topic in computational photography. Such techniques are making it possible to create virtual images of a person or place from new viewpoints and in any form of illumination.

The first photograph (Figure 1) sits in an oxygen-free protective case at the Harry Ransom Center at the University of Texas at Austin. Its image is faint enough that discerning the roofs of Joseph Nicéphore Niépce's farmhouses requires some effort.

That same effort, however, reveals a clear and amazing truth: Through the technology of photography, a pattern of light crossing a French windowsill in 1826 is re-emitted thousands of miles away nearly two centuries later. Not surprisingly, this ability to provide a view into a different time and place has come to revolutionize how we communicate information, tell stories, and document history.

DIGITAL IMAGING

In digital imaging terminology, we would think of the first photograph as a 2D array of pixel values, each representing the amount of light arriving at the camera from a particular angle. If we index the pixel values based on the horizontal (θ) and vertical (ϕ) components of this angle, we can express the photograph as a 2D function, $P(\theta, \phi)$.

Modern photographic techniques have made it possible to capture considerably more information about the light in a scene. Color photography records another dimension of information: the amount of incident light for

a range of wavelengths, λ . We can thus describe a color image as a 3D function, $P(\theta, \phi, \lambda)$. Motion pictures record how the incident light changes with time t , adding another dimension to our function, which becomes $P(\theta, \phi, \lambda, t)$.

Most of the photographic imagery we view today—cinema and television—records and reproduces light across all of these dimensions, providing an even more compelling experience of scenes from other times and places.



Figure 1. The first photograph, Joseph Nicéphore Niépce, 1826.

Source: <http://palimpsest.stanford.edu/byorg/abbey/an/an26/an26-3/an26-307.html>.




Figure 2. Geometry of the plenoptic function and the reflectance field.

Many innovations in image recording technology receiving increased interest today provide ways of capturing $P(\theta, \phi, \lambda, t)$ with greater resolution, range, and fidelity. For example, high dynamic range (HDR) imaging techniques increase the range of values that are recorded within P , accurately capturing the millions-to-one range brightness from a dimly lit interior to staring into the bright sun. Panoramic and omnidirectional photography¹ increase the range of θ and ϕ that are recorded, in some cases capturing the entire sphere of incident light. Large-format photography systems such as Graham Flint's Gigapxl Project dramatically increase angular resolution, achieving tens of thousands of independent samples across θ and ϕ . Multispectral and hyperspectral photography, respectively, increase the resolution and range recorded in λ , and high-speed photography increases the resolution in t .

As we consider what the future of digital photography might bring, we might ask what other dimensions of information about a scene we could consider capturing. In 1991, E.H. Adelson and J.R. Bergen² presented a mathematical description of the totality of light within a scene as the seven-dimensional *plenoptic function* P :

$$P = P(x, y, z, \theta, \phi, \lambda, t)$$

The last four dimensions of this function are familiar, representing the azimuth, inclination, wavelength, and time of the incident ray of light in question, as Figure 2 shows. The additional three dimensions (x, y, z) denote the 3D position in space at which the incident light is being sensed. In the case of the first photograph, (x, y, z) would be fixed at the optical center of Niépce's rudimentary camera.

Tantalizingly, this simple function contains within it every photograph that could ever have possibly been taken. If it were possible to query the plenoptic function with the appropriate values of $x, y, z, \theta, \phi, \lambda$, and t , it would be possible to construct brilliant color images (or movies) of any event in history.

VIRTUAL CONTROL OF THE VIEWPOINT

Recording variation in the spatial dimensions of the plenoptic function allows scenes to be recorded three-dimensionally. Stereo photography captures two image samples in (x, y, z) corresponding to the positions of a left and right camera. Since this matches the form of the imagery that our binocular visual system senses, it can reproduce the sensation of a 3D scene for a particular point of view.

QuickTime VR panoramas allow virtual control the field of view by panning and zooming in a panoramic image, exploring different ranges of θ and ϕ but staying at the same viewpoint (x, y, z) . In 1995, Leonard McMillan and Gary Bishop presented a system for interpolating between panoramic images taken from different viewpoints to create the appearance of virtual navigation through a scene.³

Light field photography uses an array of cameras to record a scene from a planar 2D array of viewpoints distributed across x and y for a fixed z .⁴ If the cameras are clear of obstructions, we can infer values of the plenoptic function at other z positions by following these rays forward or backward in the direction (θ, ϕ) to where the light was sensed at the x - y plane.⁵

By recording such data, it becomes possible to virtually construct any desired viewpoint of a scene *after* it has been filmed. As such techniques mature—and if display technology makes similar advances⁶—future observers may be able to enjoy 3D interactive views of sports matches, family gatherings, and tours of the great sites of the world.

This rich line of photographic advances, each recording a greater dimensionality of the light within a scene, has progressively brought a greater sense of “being there” in a scene or that the photographed subject is in some way present within our own environment. Nonetheless, a photographic image—even a 3D, omnidirectional, high-resolution motion picture—shows a scene only the way it was at a particular time, the way things happened, in the illumination that happened to be there.

If it were somehow possible to record the scene itself, rather than just the light it happened to reflect, there would be a potential to create an even more complete and compelling record of people, places, and events. In this article, we focus on a particular part of this problem: Is there a way to record a subject or scene so that we can later produce images not just from every possible viewpoint, but in any possible illumination?

If it were somehow possible to record the scene itself, rather than just the light it happened to reflect, there would be a potential to create an even more complete and compelling record of people, places, and events. In this article, we focus on a particular part of this problem: Is there a way to record a subject or scene so that we can later produce images not just from every possible viewpoint, but in any possible illumination?

VIRTUAL CONTROL OF ILLUMINATION

Recording a scene so that its illumination can be created later has potential applications in a variety of disciplines. With virtual relighting, an architect could record a building to visualize it under a proposed nighttime lighting design or capture a living room's response

to light to see how its appearance would change if an apartment complex were built across the street. An archaeologist could compare digital records of artifacts as if they were sitting next to each other in the same illumination, a process difficult to perform with standard photographs.

Some of the most interesting virtual relighting applications are in the area of filmmaking. A small crew could capture a real-world location to use it as a virtual film set, and the production's cinematographer could light the set virtually. Then, actors filmed in a studio could be composited into the virtual set, their lighting being virtually created to make them appear to reflect the light of the environment and further shaped by the cinematographer for dramatic effect.

Virtually changing the lighting in a scene is a more complex problem than virtually changing the viewpoint. Even recording the entire 7D plenoptic function for a scene reveals only how to construct images of the scene in its original illumination. If we consider a living room with three lamps, turning on any one of the lamps will give rise to a different plenoptic function with remarkably different highlights, shading, shadows, and reflections within the room.

To achieve virtual control of the illumination in the room—that is, to generate images of the room as it would appear with any set of light sources positioned anywhere in the room—it seems that we would need to capture every possible plenoptic function the room could exhibit. Since capturing just one plenoptic function requires complex instrumentation and a significant amount of data, this is a challenging proposition.

At this point, a fact from physics comes to our aid: Light is additive. This means that if we want to show the appearance of the living room from a particular viewpoint under any of the 2^3 combinations of the three lamps being on, we only need to take three basis images, with one lamp on in each image. To generate an image with any set of the lights on, we can simply add together the pixel values of the corresponding single-light images, producing a result accurate to within the measurement error of the sensor.

Going further, we can predict the scene's appearance if one of the lights were twice as bright by adding that light's image twice to the sum. We can even simulate the effect of changing the color of one of the light sources by scaling the color channels of the basis images before adding them together. In 1992, Paul Haeberli used this technique to show a scene with blue light coming from one direction and red light coming from another by tinting two images lit by white light sources and then adding them together.⁷ This is good news for being able to sim-

ulate arbitrary lighting conditions in a scene, since we only need to capture the scene in a set of basis lighting conditions that span the space of the lighting conditions of interest. In 1995, Julie Dorsey and colleagues suggested this technique in the context of computer-rendered imagery to efficiently visualize potential theatrical lighting designs.⁸

We should now ask which set of basis lighting conditions will allow us to simulate *any* possible pattern of illumination in the scene. Since we would like to place lights at any (x, y, z) position, we might suppose that our basis set of recorded plenoptic functions should consist of a set of images of the scene, each with an omnidirectional light source (such as a tiny lightbulb) at a different (x, y, z) position, densely covering the space. However, this lighting basis would not allow us to simulate how a directed spotlight would illuminate the scene, since no summation of omnidirectional light source images can produce a pattern of light aimed in a particular direction.

Thus, our basis must include variance with respect to lighting direction as well, and we should choose our basis light sources to be single rays of light, originating at (x, y, z) and emitting in direction (θ, ϕ) . From such a set of rays, we can construct any possible lighting environment consisting of any number of artificial lights, natural illumination from the sun and sky, or indirect light from other surfaces around the scene.

We can see that the space of incident illumination begins to resemble the 7D plenoptic function itself. Indeed, to be general, we must also consider the incident ray of light's wavelength λ and its time of emission t . Considering these seven dimensions together, for both an incident ray of light L_i and the plenoptic function of radiant light L_r , it produces, yields a 14D version of the *reflectance field* R :⁹

$$R = R(L_i; L_r) = R(x_i, y_i, z_i, \theta_i, \phi_i, t_i, \lambda_i; x_r, y_r, z_r, \theta_r, \phi_r, t_r, \lambda_r)$$

We can imagine recording this function using a laser pointer emitting light at a controllable wavelength λ_i and a tiny light sensor with a narrow opening recording light at a controllable wavelength λ_r . If we place the laser pointer in our scene such that its opening is at (x_i, y_i, z_i) aimed toward (θ_i, ϕ_i) , emitting a short pulse of light at time t_i , then R would be the amount of energy recorded by our sensor placed at (x_r, y_r, z_r) receiving light from direction (θ_r, ϕ_r) at time t_r . If such a data set could be captured for a real scene, it would be possible to perform both of the principal aspects of cinematography—making camera and lighting choices—as a virtual postproduction process. However, capturing this function

Some of the most interesting virtual relighting applications are in the area of filmmaking.



Figure 3. The Light Stage 5 performance relighting process. (a) A sphere of 156 white LED light sources surrounds the actor as a microcontroller changes the lighting direction several thousand times per second. A high-speed digital video camera captures the actor's appearance under each lighting direction. (b) 156 consecutive high-speed frames of the actor's performance captured in just 1/24th of a second, each showing a different basis lighting direction. (c) The actor's performance synthetically illuminated by two real-world lighting environments (Grace Cathedral, top; the Uffizi Gallery, middle) and a synthetic lighting environment (below) formed as linear combinations of the basis lighting directions.

exhaustively even for a limited region of space is much more daunting than capturing just a single plenoptic function, and conjures images of robotic armatures moving lasers and sensors throughout a scene, recording terabytes upon terabytes of information.

Fortunately, a few reasonable simplifications make the reflectance field more tractable. One is that light projected into the scene at wavelength λ_i usually stays the same wavelength as it reflects and scatters within the scene—a fact that is generally true in the absence of fluorescent materials. Another simplification is that light emitted at time t_i will arrive at essentially the same time at the sensor—a reasonable approximation for human-scale scenes and timescales.

Each of these simplifications eliminates one dimension of information (λ_i and t_i , respectively) from the function. A curious fact about this function is that if the positions of the light emitter and the light sensor were transposed, the same amount of energy would be recorded, regardless of the scene's complexity—a phenomenon known as *Helmholtz reciprocity*. However, while this fact theoretically halves the amount of information present in R , it does not reduce its dimensionality.

Projects to date that capture scenes so that both lighting and viewpoint can be chosen virtually make additional assumptions so that the acquisition process is tractable. One such project undertaken by our laboratory films human performances so that their lighting can be designed in postproduction, and another records a relightable model of an entire outdoor environment.

Each project makes different assumptions to achieve a useful result for its subject.

RELIGHTING A HUMAN PERFORMANCE

The first example is a technique for relighting human performances.¹⁰ The Light Stage 5 apparatus shown in Figure 3 is a 2-meter sphere of 156 white LED light sources that surround an actor. The LED lights are controlled by a microcontroller that can change the lighting direction thousands of times per second, fast enough that the illumination appears as a fluttering sphere of light rather than sequential lighting directions. Filming the actor with a synchronized high-speed video camera yields a stream of images under the repeating sequence of 156 lighting directions, with each complete sequence taking just 1/24th of a second to capture.

The data captured with this setup spans six dimensions of the 14D reflectance field. There is just one camera viewpoint, fixing (x, y, z) , but each image in the sequence spans a field of view in (θ, ϕ) . If we consider the lights to be distant from the face, we can consider each light to represent a unique lighting direction in (θ, ϕ) . However, since the lights project over the whole face (rather than projecting single points with single rays), there is no spatial variation in the illumination, eliminating variability across (x, y, z) .

We further assume there is no fluorescence (eliminating the need to consider λ_i independently of λ_r , which allows us to light the subject simultaneously with all wavelengths of visible light) and that light travels instantaneously (eliminating the need to consider t_i independently of t_r). Variability of the reflectance field in time t , is effectively captured 24 times per second—the frame rate of motion pictures—and variability with wavelength λ_r is captured in the image sensor's red, green, and blue color channels. Thus, the dimensions in which variability in the reflectance field is recorded can be written:

$$R' = R'(\theta, \phi; \theta, \phi, t, \lambda_r)$$

From this data, we can show how the actor's performance would appear under any number of lighting conditions by summing scaled and tinted versions of the



Figure 4. A frame from a performance by actor Chris Carmack, directed by Randal Kleiser, virtually rendered with nine different lighting configurations.

basis images. Each of the images in Figure 4 was formed by adding together two or more of the basis lighting directions for a particular moment of an actor's performance, with each basis light chosen to represent an interesting key, fill, or rim light direction. The significant changes in the actor's appearance across this range of illumination give evidence of the rich information contained in a relightable image data set. In Figure 3c, a different actor's performance is illuminated by adding together all 156 images, with the color channels of each image scaled according to the color and intensity of the incident illumination of three different lighting environments. The first two results show the actor as she would appear standing inside Grace Cathedral in San Francisco or in the Uffizi Gallery in Florence, Italy, based on actual illumination recorded at these locations.

While this relighting technique produces realistic results for light from a distant environment, it does not simulate how the actor would appear in dappled light or in partial shadow. However, such spatially varying illumination is frequently used in cinematography for dramatic effect: placing uneasy focus on a character by casting brighter light across the eyes, or showing a character lit through Venetian blinds as a foreboding visual in film noir. To perform such relighting by recombining captured data, we would need to capture spatial light-




Figure 5. Simulating spatially varying illumination on an actor's performance. (a) The actor is lit by a structured illumination pattern from a high-speed video projector. (b) A geometric model is derived for a frame of the performance. (c) The actor is virtually illuminated as if in a room with sunlight streaming through Venetian blinds. (d) The actor is virtually illuminated as if in a cathedral with light projected from a stained glass window.

ing variation across (x, y, z) , returning us to the impractical situation of needing to capture the effect of incident illumination one ray of light at a time.

Fortunately, we can approximate the effect of spatially varying illumination by virtually lighting each pixel of the face as if in a different lighting environment based on where it lies within the desired virtual field of illumination. This requires knowing the 3D position of each pixel in each frame of the performance, which can be determined by performing a 3D scan of the actor's geometry for each frame.¹¹

As Figure 5 illustrates, we have done this using a high-speed video projector from Fakespace Labs to project a series of structured illumination patterns onto the actor, interleaved with the light stage's basis lighting conditions. We use the appearance of each pixel under the series of stripe patterns to triangulate its 3D position and produce a surface model of each frame of the performance, shown in Figure 5b. From this geometry, we illuminate each pixel as if it were receiving different environmental illumination. Figure 5c shows a frame of a performance synthetically illuminated as if in a room with the sun streaming through Venetian blinds, and Figure 5d shows the performance illuminated as if in a cathedral with projected light from a stained glass window.

Recent work in our laboratory has endeavored to advance beyond two remaining limitations of this relighting process. One advance has been to apply the relighting process to full-body performances, for which we have constructed Light Stage 6, the 8-meter dome seen in Figure 6a, which evenly focuses LED illumination over a person standing at the center.

The other advance has been to include virtual control of the camera viewpoint, recording variation across the dimensions of (x_r, y_r, z_r) . We have achieved this result for the special case of *cyclic* human motion such as walking and running by filming the actor on a rotating treadmill with a vertical 1D array of high-speed cameras.¹² As the actor walks and rotates, the cameras effectively record the locomotion cycle for a circle of horizontally spaced viewpoints as well. Using a combination of view interpolation and light field rendering, the actor's motion can be rendered from any (x_r, y_r, z_r) point of view as well as illuminated by any virtual lighting environment. Figure 6b shows a frame from an animation that used this technique to virtually render

several instances of a running subject in the courtyard at the Doge's Palace in Venice.

RELIGHTING A LARGE ENVIRONMENT

The light stage technique is a useful process for capturing relightable imagery of people, and parts of this process have been used to create relightable digital stunt doubles of the main characters in motion pictures such as *Spider-Man 2*, *King Kong*, and *Superman Returns*. Similar techniques have been used to capture relightable representations of human-scale objects such as archaeological artifacts as well. However, a different technique must be used to record a relightable representation of an outdoor environment, since constructing a light stage large enough to contain such a scene would be impractical. Recently, we addressed this problem in the context of capturing a relightable model of the Parthenon in Athens.¹³

To solve this problem, we photographed the Parthenon from different positions (x_r, y_r, z_r) in different natural lighting conditions (θ, ϕ) , each providing samples of the scene's radiant plenoptic function L_r under that lighting. Each time we took such a photograph, we simultaneously captured an image of the incident illumination L_i using the light probe device shown in Figure 7a, which

records the full dynamic range of the sky's illumination across (θ, ϕ) .

We then used a time-of-flight laser scanner to measure the 3D surface model shown in Figure 7b, and we assumed as an initial estimate that the Parthenon model's surface reflectance properties were a neutral diffuse color. We then used the *Arnold* physically based lighting simulation program to render virtual images of the model as it would appear from the viewpoint in our photograph L_r , using the captured illumination L_i as the virtual light in each rendering. Since the scene geometry, camera viewpoint, and incident illumination were matched in each photograph-rendering pair, pixel value discrepancies were taken to be errors in the virtual model's reflectance properties as compared to those of the real scene. We used these discrepancies to iteratively adjust the model's surface reflectance properties—brightening the surfaces that appeared too dark in the renderings and darkening those that appeared too bright.

After a few iterations of such adjustments, this produced a virtual 3D model of the Parthenon which, when illuminated with any of the captured lighting environments L_i , accurately reproduced the real Parthenon's appearance under that light as recorded in the photograph L_r .

Capturing a relightable representation of the Parthenon in this way required making significant assumptions about its reflectance field. Specifically, we assumed that light reflection happened at solid surface bound-




Figure 6. The Light Stage 6 performance relighting system. (a) A running subject is recorded in Light Stage 6, an 8-meter sphere that captures dynamic subjects under a wide array of viewing and lighting directions. (b) Several instances of the runner are rendered into a photographed environment, each with its viewpoint and illumination chosen to be consistent with the virtual scene.



Figure 7. Capturing the lighting and geometry of the Parthenon. (a) A light probe device records incident lighting conditions near the base of the temple. (b) A 3D surface model of the scene geometry obtained using 3D laser scanning.



Figure 8. *The Parthenon virtually reilluminated. (a) A virtual model of the Parthenon rendered as it would appear under natural illumination (inset) recorded in Marina del Rey, California. (b) The reilluminated model rendered with a nighttime lighting design.*

aries (that is, at the scanned 3D geometry) and that the surface reflectance effects were well-characterized by relatively diffuse (rather than shiny) materials. Nonetheless, the model realistically depicts the Parthenon as it would appear under essentially any new lighting conditions.

Figure 8a shows a virtual rendering of the Parthenon lit by a dramatic sky photographed on the roof of our institute in Marina del Rey, California. The inset shows the high dynamic range illumination data set, shot as an HDR fisheye image. Figure 8b was created by placing virtual spotlights in the scene, visualizing a potential nighttime lighting design.

LOOKING TO THE FUTURE

Recording how scenes transform incident illumination into radiant light is an active topic in computational photography. Much of the work focuses on designing more advanced basis illumination conditions to acquire

the reflectance field more efficiently and completely.

Yoav Schechner and colleagues explore how using basis illumination conditions with more than one light element on at a time can record reflectance field data with an improved signal-to-noise ratio.¹⁴ Pieter Peers and Philip Dutré analyze how a scene reflects patterns of wavelet noise from a video projector, and by assuming that the scene's reflectance field is not noisy, estimate the scene reflectance with a substantially reduced number of lighting conditions.¹⁵ Pradeep Sen and colleagues use the reciprocity principle of light transport to capture reflectance interchangeably with both projectors and cameras, and they analyze the captured data while it is being recorded to adaptively focus on the scene's most significant reflectance information.¹⁶

Some of the most interesting future work might endeavor to analyze data from reflectance fields—recorded from light generated and sensed at the periphery of the scene—to infer how each infinitesimal region *within* the scene reflects and scatters illumination. Such a scattering field will describe the scene's reflectance more compactly as a volumetric data set of local light interaction, maintaining generality, in a manner that is closer to representing the matter of the scene itself.

There is something new and exciting in the ability to record reilluminable imagery of people, objects, and environments. Instead of having just photographs—views into the past of how a subject once was—we have a representation which we can relight and reinterpret in a new creative vision, and taking a step closer to preserving the subject itself. While the importance of these new techniques can never rival the original impact of photography, the quality of seeing a person's face, virtually reflecting the light of a place she has never been to, must channel a small and distant echo of the magic in the first photograph. ■

Animated movies from the projects shown in this article can be found at <http://gl.ict.usc.edu/research/LS5/> and www.debevec.org/Parthenon.

Acknowledgments

These projects have been done in collaboration with Mark Bolas, Marc Brownlow, Charles-Felix Chabert, Lora Chen, Aimee Dozois, Per Einarsson, Brian Emerson,

Marcos Fajardo, Andrew Gardner, Tim Hawkins, Andrew Jones, Bruce Lamond, Therese Lundgren, Alex Ma, Ian McDowall, Maya Martinez, Philippe Martinez, Tom Pereira, Diane Piepol, Charis Poullis, John Shipley, Jessi Stumpf, Sebastian Sylwan, Chris Tchou, Jonas Unger, Andreas Wenger, and Nathaniel Yun, and have been supported by USA RDECOM, the University of Southern California Office of the Provost, and TOPPAN Printing Co., Ltd. The information presented here does not necessarily reflect the position or policy of the sponsors and no official endorsement should be inferred.

References

1. S.K. Nayar, "Omnidirectional Vision," *Proc. 8th Int'l Symp. Robotics Research*, Springer, 1998, pp. 195-202; www1.columbia.edu/CAVE/publications/pdfs/Nayar_ISRR97.pdf.
2. E.H. Adelson and J.R. Bergen, "The Plenoptic Function and the Elements of Early Vision," M. Landy and J.A. Movshon, eds., *Computational Models of Visual Processing*, MIT Press, 1991, pp. 3-20.
3. L. McMillan and G. Bishop, "Plenoptic Modeling: An Image-Based Rendering System," *Proc. 22nd Ann. Conf. Computer Graphics and Interactive Techniques*, ACM Press, 1995, pp. 39-46.
4. M. Levoy, "Light Fields and Computational Imaging," *Computer*, Aug. 2006, pp. 46-55.
5. M. Levoy and P. Hanrahan, "Light Field Rendering," *Proc. 23rd Ann. Conf. Computer Graphics and Interactive Techniques*, ACM Press, 1996, pp. 31-42.
6. O. Bimber, "The Ultimate Display—What Will It Be?" *Computer*, Aug. 2005, pp. 28-29.
7. P. Haeberli, "Synthetic Lighting for Photography," *Grafica Obscura*, Jan. 1992; www.graficaobscura.com/.
8. J. Dorsey, J. Arvo, and D. Greenberg, "Interactive Design of Complex Time-Dependent Lighting," *IEEE Computer Graphics and Applications*, vol. 15, no. 2, 1995, pp. 26-36.
9. P. Debevec et al., "Acquiring the Reflectance Field of a Human Face," *Proc. 27th Ann. Conf. Computer Graphics and Interactive Techniques*, ACM Press, 2000, pp. 145-156.
10. A. Wenger et al., "Performance Relighting and Reflectance Transformation with Time-Multiplexed Illumination," *ACM Trans. Graphics*, vol. 24, no. 3, 2005, pp. 756-764.
11. A. Jones et al. "Performance Geometry Capture for Spatially Varying Relighting," tech. sketch, *Proc. 32nd Int'l Conf. Computer Graphics and Interactive Techniques*, 2005; http://gl.ict.usc.edu/research/SpatialRelighting/spatial_sketch.pdf.
12. P. Einarsson et al., "Relighting Human Locomotion with Flowed Reflectance Fields," *Proc. 17th Eurographics Symp. Rendering*, Springer-Verlag, 2006.
13. P. Debevec et al., *Estimating Surface Reflectance Properties of a Complex Scene under Captured Natural Illumination*, tech. report ICT-TR-06-2004, USC ICT, 2004; <http://gl.ict.usc.edu/research/reflectance/Parth-ICT-TR-06.2004.pdf>.
14. Y.Y. Schechner, S.K. Nayar, and P. Belhumeur, "A Theory of Multiplexed Illumination," *Proc. 9th Int'l Conf. Computer Vision*, IEEE CS Press, 2003, pp. 808-815.
15. P. Peers and P. Dutré, "Inferring Reflectance Functions from Wavelet Noise," *Proc. Eurographics Symp. Rendering Techniques*, Eurographics Assoc., 2005, pp. 173-182.
16. P. Sen et al., "Dual Photography," *ACM Trans. Graphics*, vol. 24, no. 3, 2005, pp. 745-755.

Paul Debevec is a research associate professor in the University of Southern California Computer Science Department and the executive producer of graphics research at the USC Centers for Creative Technologies. His research interests focus on computer vision and computer graphics for photoreal scene acquisition and rendering. Debevec received a PhD in computer science from the University of California, Berkeley. He is a member of ACM Siggraph and the Visual Effects Society. See www.debevec.org for additional information.

Get access

to individual IEEE Computer Society documents online.

More than 100,000 articles and conference papers available!

\$9US per article for members

\$19US for nonmembers

www.computer.org/publications/dlib

IEEE
 computer
 society
 60th anniversary

