

Statistics of Real-World Illumination

CVPR Paper Number 613

Abstract

While research in computer vision often assumes simple illumination models, real-world illumination is highly complex, consisting of reflected light from every direction as well as distributed and localized primary light sources. One can capture the illumination incident at a point in the real world from every direction photographically in the form of a spherical illumination map. This paper illustrates, through analysis of photographically-acquired, high dynamic range illumination maps, that real-world illumination shares many of the statistical properties of natural images. In particular, the marginal and joint wavelet coefficient distributions, directional derivative distributions, and harmonic spectra of illumination maps resemble those documented in the extensive literature on natural image statistics. However, illumination maps differ from standard natural images in that illumination maps are statistically non-stationary and may contain localized light sources which dominate their power spectra. Our work provides a foundation for statistical models of real-world illumination that may facilitate robust estimation of shape, reflectance, and illumination from images.

1. Introduction

In an effort to design vision systems which function robustly outside the laboratory, researchers have recently devoted a great deal of effort to capturing the statistics of the world. They have studied statistics of natural optical images [10, 7, 21, 19], range images [9] and Gestalt grouping factors [15]. These statistics have proven useful for explaining the architecture of biological visual systems [7, 13], developing efficient image coding schemes [21], denoising images [18], setting parameters for grouping algorithms [15], and more generally providing models for probabilistic or Bayesian modeling.

The focus of this paper is to understand the statistics of natural illumination. In traditional computer vision, illumination often refers to luminous sources alone. For example, shape from shading calculations typically assume a single point light source at infinity or a uniform hemispherical source. In the real world, however, illumination is much more complicated. Consider the rendering equation in computer graphics. The total reflected radiance of a

surface patch in the view direction (θ_r, ϕ_r) is given by:

$$\int_{\phi_i=0}^{2\pi} \int_{\theta_i=0}^{\pi/2} f(\theta_i, \phi_i; \theta_r, \phi_r) I(\theta_i, \phi_i) \cos \theta_i \sin \theta_i d\theta_i d\phi_i \quad (1)$$

where $I(\theta_i, \phi_i)$ is the radiance of illumination incident from direction (θ_i, ϕ_i) . In other words, the appearance of a surface patch is affected by illumination incident from *all* directions. Illumination includes not only direct illumination from the sun, sky, or indoor lights, but also indirect illumination, by light reflected from other surfaces in the environment.

One can measure the illumination $I(\theta_i, \phi_i)$ incident from every direction at a particular point in the real world using a camera whose optical center is located at the point of interest. By combining photographs representing illumination from every direction, one can compose a spherical illumination map describing $I(\theta_i, \phi_i)$. Such spherical images are used as environment maps in computer graphics [4]. If all sources of direct and indirect illumination are relatively distant, the illumination map changes little as the hypothetical camera moves through space.

An illumination map is a type of image. However, accurate real-world illumination maps differ from the photographs studied in the natural image statistics literature in several regards. First, illumination maps cover a much wider view angle. While previously studied photographs have typically cover a narrow view angle near the horizontal, illumination maps ideally cover the entire sphere. Second, accurate illumination maps possess a much higher dynamic range than previously analyzed photographs. In fact, they may contain localized primary light sources such as incandescent lights or the sun itself.

Researchers have found that normal photographs of indoor and outdoor scenes display a great deal of regularity, particularly in power spectra and distributions of wavelet coefficients [7, 19, 10]. We wish to determine the similarities and differences of these statistics to those of high dynamic range, photographically-acquired illumination maps. Working with two sets of illumination maps, we therefore analyze marginal and joint distributions of illumination intensity (Sections 4 and 6), distributions of directional intensity derivatives (Section 5), spherical harmonic power spectra (Section 7), and marginal and joint wavelet coefficient distributions (Section 8). When possible, we compare our results to those of Huang and Mumford [10], who

analyzed a set of over 4000 restricted-angle outdoor photographs collected and calibrated by van Hateren and van der Schaaf [24]. Like Huang and Mumford, we typically work with log pixel intensities and subtract out the mean of the log of each image to normalize for overall brightness.

We find that the statistics of illumination are surprisingly similar to those of more typical photographs. We encountered several significant differences, however, some of which are due to the marked non-stationarity of illumination statistics and to the presence of concentrated light sources which can dominate power spectral measures.

Our work provides a foundation for an understanding of illumination statistics that may serve several practical purposes. First, prior statistical information about real-world illumination is necessary for Bayesian estimation of surface reflectance properties [6]. The problem of recovering a surface BRDF from an image by inverting equation Eq. 1 is ill-posed under unknown illumination, but can be solved by exploiting such prior information. A statistical description of illumination can also improve the recovery of illumination fields from incomplete data, an important problem in computer graphics [14]. Finally, an accurate description of illumination statistics may facilitate the development of shape-from-shading algorithms which function robustly under complex, unknown illumination.

2. Datasets

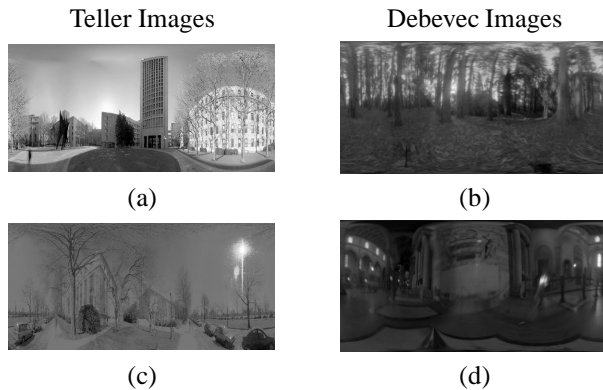


Figure 1: Examples of the illumination maps we used, shown in equal-area cylindrical projection. (a) and (c) are drawn from Teller’s data set, while (b) and (d) are drawn from Debevec’s. Pixel values have been passed through a compressive nonlinearity for display purposes.

We worked with two different data sets, each consisting of high dynamic range images that represent the radiance incident from all directions at a point in the real world. The first data set consisted of 95 illumination maps based on imagery acquired by Teller et al. [23] in the environs of the MIT campus (<http://city.lcs.mit.edu/data>). The second set

consisted of nine maps from Debevec’s Light Probe Image Gallery (<http://www.debevec.org/Probes/>) [3]. Debevec’s maps represent diverse lighting conditions from four indoor settings and five outdoor settings. Two examples from each data set are shown in Figure 1.

The images in both data sets were acquired by combining photographs at multiple exposures to obtain pixel values which are linear in luminance, using the technique of Debevec and Malik [5]. We converted them all to gray-scale images that are logarithmic in luminance. Debevec’s illumination maps, which were computed from photographs of a chrome ball, cover the entire sphere. Teller’s illumination maps were each mosaiced from multiple calibrated narrow field of view images. These mosaics cover the entire upper hemisphere as well as a band below the equator.

3 Spherical projection

Whereas image statistics have previously been analyzed on a planar domain, illumination maps are naturally defined on a sphere. We will describe our handling of this issue in each of the following sections. We found that storing the illumination map in an equal area cylindrical projection [2] facilitated certain computations. To construct this projection, one places the sphere at the center of a vertically oriented cylinder and projects each point on the spherical surface horizontally outward to the surface of the cylinder (Figure 2). One then unwraps the cylinder to obtain a rectangular map of finite extent. Regions of equal area on the sphere map to regions of equal area on the cylinder.¹ Figure 1 displays illumination maps in equal-area projection with $k = \frac{2}{\pi}$, where k is the ratio of the radius of the cylinder to the radius of the sphere.

4. Illumination Intensity Distribution

To compute the distribution of illumination intensities incident from all directions, one must take into account the solid angle corresponding to each pixel of the illumination map. For an equal area projection, this solid angle is constant, so we can estimate the distribution with a simple pixel histogram. Figure 3 shows total illumination intensity distributions for the 95 Teller images and for the 9 Debevec images. Huang and Mumford [10] noted asymmetry in their single pixel distribution due to the presence of sky in many of their images. Our distributions exhibit more striking asymmetries, partly because both the Teller and Debevec data sets contain not only sky but other localized light sources. The

¹In particular, an infinitesimal patch on the sphere at latitude θ will find itself expanded by a factor of $k \frac{1}{\cos \theta}$ in the horizontal direction and reduced by a factor of $\cos \theta$ in the vertical direction. Because the product of these two factors is a constant k , this projection preserves areas, even though it heavily distorts angles near the poles.




Figure 2: To produce the equal area cylindrical projection of a spherical map, one projects each point on the surface of the sphere horizontally outward onto the cylinder, then unwraps the cylinder to obtain a rectangular map.

distribution for the Teller set is particularly asymmetric due to the presence of the sun in many images and to undersaturation in the imaging system at very low light intensities. Our distributions are also much noisier than Mumford’s because we averaged over fewer images. Pixel intensity distributions vary much more from image to image than wavelet statistics or power spectra (Sections 8 and 7).

Following Huang and Mumford [10], we computed the standard deviation σ , skewness S , kurtosis κ , and differential entropy \mathcal{H} .² For the Teller images, we found $\sigma = 1.04$, $S = -0.02$, $\kappa = 4.04$, and $\mathcal{H} = 2.06$. For the Debevec images, we have $\sigma = 1.32$, $S = 0.36$, $\kappa = 12.49$, and $\mathcal{H} = 2.21$ for the Debevec images. Huang and Mumford found $\sigma = 0.79$, $S = 0.22$, $\kappa = 4.56$, and $\mathcal{H} = 1.66$. Hence our illumination maps have a higher variance and entropy, and a much larger kurtosis. This difference is due to the higher dynamic range of our images and to the inclusion of bright localized light sources, including the sun.




Figure 3: Illumination intensity distribution: log histogram of $\ln(I(\theta_i, \phi_i)) - \text{mean}(\ln(I))$. Left: statistics for 95 Teller images. Right: statistics for 9 Debevec images.

²These distributions have mean 0, because we subtract out the mean log value before processing.

4.1. Non-stationarity

Illumination statistics vary with direction. Figure 4(a) and (b) show mean luminance as a function of elevation. As expected, illumination generally increases with elevation. Interestingly, the mean intensity reaches a local minimum at the horizontal view direction because both data sets contain illumination maps in which the ground reflects a significant amount of light from above, while visible surfaces in the horizontal direction are shadowed (e.g., Figure 1b). Panels (c) and (d) of Figure 4 each show two histograms at different ranges of elevations. The histograms for higher view directions have not only a larger mean but heavier positive tails, reflecting the larger probability of bright localized sources at higher elevations. Lack of statistical stationarity is seldom reported in the natural image literature due to the limited field of view of the images analyzed. It has, however, been observed in range data [9].

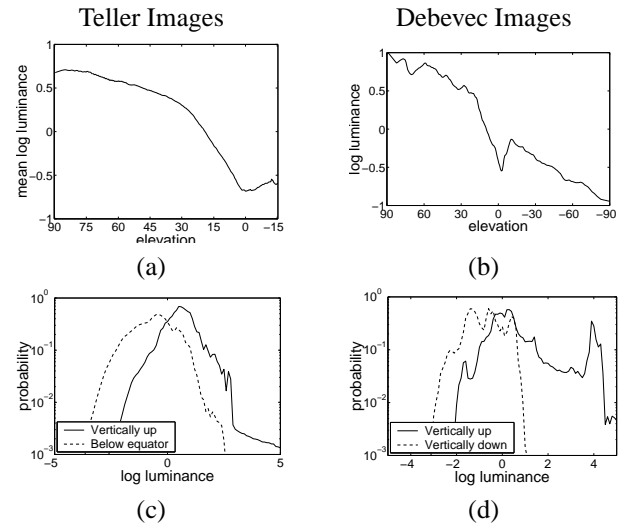


Figure 4: Dependence of illumination on elevation. (a) and (b) show mean luminance as a function of elevation for the Teller and Debevec images respectively. (c) and (d) each show two histograms, one for a range of angles near the vertically upward direction, the other for a range of angles at a lower elevation.

5. Derivative Statistics

Following Huang and Mumford, we approximate horizontal derivatives as differences between horizontally adjacent samples and compute their marginal distribution. We define the horizontal direction in the global coordinate frame, such that horizontal derivatives correspond to differences along lines of latitude. Because lines of latitude differ in length, we define horizontally adjacent positions as being separated by a fixed distance on the sphere. This distance is chosen

such that the equator is divided into 512 parts.




Figure 5: Solid lines indicate distribution of horizontal derivatives for log luminance illumination maps. Dashed lines are generalized Laplacian fits.

Figure 5 shows the resulting marginal distributions of horizontal derivatives for the two datasets. Both distributions, like that of Huang and Mumford, are symmetric with very high kurtoses. We found $\kappa = 15.9$ for the Teller images and $\kappa = 341.86$ for the Debevec images, while Huang and Mumford found $\kappa = 17.4$.³ Interestingly, the variance of our distributions ($\sigma = 0.45$ for the Teller images and $\sigma = 0.26$ for the Debevec images) is also similar to that reported by Huang and Mumford ($\sigma = 0.26$), even though the angular distance between their adjacent samples is approximately seven times smaller than ours.

We fit generalized Laplacian distributions of the form $\mathcal{P}_x(x) \propto \exp(-|x/s|^\alpha)$ to both histograms using a maximum-likelihood criterion, obtaining $\alpha = 0.53$ and $s = 0.055$ for the Teller images and $\alpha = 0.56$ and $s = 0.034$ for the Debevec images (Figure 5). While these parameters are similar to those of Huang and Mumford, who found $\alpha = 0.55$, our distributions are not as well modeled by generalized Laplacians as theirs. The shapes of our distributions differ somewhat from each other and from that of Huang and Mumford, possibly because of the different distribution of luminous sources in each.

6. Joint Distribution of Illumination from Two Adjacent Directions

Again following Huang and Mumford, we computed the joint distribution of log intensities at horizontally adjacent positions in the illumination maps. Specifically, we compute the joint distribution of p_1 and p_2 , where p_1 and p_2 represent log luminances at positions on the sphere which are horizontally adjacent as defined in Section 5. Figure 6 shows a contour plot of the resulting distribution over all of Teller’s illumination maps. The distribution has a shape similar to that of Huang and Mumford, who also found some asymmetry between the upper right and lower left

³As Huang and Mumford point out, the computed kurtoses are very sensitive to outliers; this is particularly true for the Debevec data set, which consists of a small number of images.




Figure 6: Joint histogram of horizontally adjacent positions in Teller’s illumination maps.

quadrants. In our case, the increased extent of the joint distribution in the upper right quadrant compared to the lower left reflects the asymmetry of the marginal distribution illustrated in Figure 3.

In agreement with Huang and Mumford, we found that while p_1 and p_2 are highly correlated, $p_1 + p_2$ and $p_1 - p_2$ are more nearly independent. In particular, the mutual information of p_1 and p_2 is 2.41 bits, while that of $p_1 + p_2$ and $p_1 - p_2$ is only .103 bits. Hence, the percentage difference between the luminance incident from two horizontally adjacent spatial directions is roughly independent of the mean luminance from those two directions.

7. Spherical harmonic power spectra

Much early work on natural image statistics focused on the regularity of power spectra. A number of authors [7, 19] have observed that two-dimensional power spectra of natural images typically fall off as $1/f^{2+\eta}$, where f represents the modulus of the frequency and η is a small constant which varies from scene to scene.

The natural equivalent of the Fourier transform on the sphere is a spherical harmonic decomposition. The spherical harmonics form a countable orthonormal basis for square integrable functions on the sphere. Associated with each basis function is an order L , a nonnegative integer analogous to frequency. The $2L + 1$ spherical harmonics of order L span a space that is closed under rotation [11].

Just as planar white noise has a flat two-dimensional power spectrum, white noise on the sphere produces equal power in every spherical harmonic. Similarly, if the regularities observed in the natural image statistics literature carry over to spherical illumination maps, the average power of the spherical harmonics at order L will fall off as $1/L^{2+\eta}$.

We computed spherical harmonic coefficients up to order 256 for the illumination maps in both data sets using the formulas given by Inui [11]. We obtained average power at each order L as the mean of squares of the coefficients at that order. Teller’s data lacks information about the lowest portion of the illumination hemisphere. We applied a smooth spatial window to these illumination maps before transforming them to the spherical harmonic domain.




Figure 7: Spherical harmonic power spectra (solid lines) of illumination maps (a), (b), (c), and (d) in Figure 1, with pixel value proportional to log luminance. The dotted lines of slope -2 correspond to power spectra of the form k/L^2 .

Figure 7 shows the relationship between average power and harmonic order for the four illumination maps of Figure 1, when pixel value is proportional to log luminance. All four images have power spectra which lie close to a straight line of slope -2 on log-log axes, corresponding to a power spectrum of the form k/L^2 . The great majority of images in both data sets exhibit similar behavior.

We obtain very different results for the same illuminations when we compute power spectra for illumination maps whose pixel values are linear in luminance. Illumination maps such as those of Figure 1a and b, which lack concentrated primary light sources, have spherical harmonic spectra that are well approximated as $k/L^{2+\eta}$ for small η . On the other hand, illumination maps that contain intense, localized light sources have smooth power spectra which remain flat at low frequencies before falling off sharply at higher frequencies. The illuminations of Figure 1c and d both display this behavior; the power spectrum of a linear luminance version of Figure 1c is shown in Figure 7. In these images, one or a few luminous sources, such as the sun or incandescent lights, dominate the power spectrum. Because these light sources approximate point sources, their spectrum is flat at low frequencies. If one clips the brightest pixel values in these images, the power spectra return to the familiar $k/L^{2+\eta}$ form (Figure ??).

Previous work on natural images has reported $1/f^{2+\eta}$ power spectra whether pixel values are linear or logarithmic in luminance [19]. These results on linear luminance images differ from ours because most previous researchers have avoided photographs of point-like luminous sources

and have used cameras of limited dynamic range, such that a few maximum intensity pixels could not dominate the image power spectra. A natural illumination map may be dominated by light sources occupying a small spatial area. Once the relative strength of such sources is reduced through clipping or a logarithmic transformation, illumination maps have power spectra similar to those of typical photographs.




Figure 8: Left, the spherical harmonic power spectrum of illumination map (c) with pixel value linear in luminance. Right, the corresponding spectrum after the pixel values corresponding to the sun have been clipped to a luminance value only slightly greater than that of the sky.

8. Wavelet Statistics

Perhaps the most powerful characterizations of natural images in the current literature are in the wavelet domain. Distributions of wavelet coefficients at various scales and orientations share a great deal of structure from image to image, as do joint distributions of wavelet coefficients at different scales, orientations, or spatial positions. A number of authors have used properties of these distributions for image denoising [18, 22], texture characterization [8, 17], or reflectance classification [6].

Prior analysis of natural images and textures has assumed that the data is defined on a planar domain. One could use spherical wavelets [20] to analyze the statistics of spherical illumination maps. In order to better compare our results with those of the natural image statistics literature, however, we elected to use planar wavelets. In particular, spherical wavelets lack the natural sense of orientation of planar wavelets. We experimented with several projections of the sphere to the plane for the purpose of wavelet analysis and obtained similar results. The plots shown in this section are based on equal area cylindrical projections, with $k = \frac{2}{\pi}$, of spherical log-luminance illumination maps.

Figure 9a shows marginal distributions of horizontally oriented Haar wavelet coefficients at three successive scales, together with maximum likelihood generalized Laplacian fits. Each marginal distribution is highly kurtotic and fits the generalized Laplacian closely. The marginal distributions increase in variance at successively coarser scales. Figure 9b shows the corresponding distributions for vertically oriented wavelets. The generalized Laplacian




Figure 9: Distributions of Haar wavelet coefficients at successive scales (thick lines), along with generalized Laplacian fits (thin lines). Left, horizontal bands. Right, vertical bands. Parameters of fits are: horizontal, finest scale, $\alpha = 0.58$, $s = 0.043$ second finest scale, $\alpha = 0.52$, $s = 0.056$, third finest scale, $\alpha = 0.5036$, $s = 0.1007$ vertical, finest scale, $\alpha = 0.54$, $s = 0.043$ second finest scale, $\alpha = 0.48$, $s = 0.052$ third finest scale, $\alpha = 0.43$, $s = 0.067$.

fit is poorest at the finest scale; this particular distribution corresponds approximately to the distribution of horizontal pixel-wise differences in Figure 5a. The poor fit may be due to noise at the finest scale.

Figure 10 shows contour plots of joint distributions of wavelet coefficients with various relationships. In these plots, “horizontal component”, “vertical component”, and “diagonal component” refer to the wavelet coefficients of different orientations at a given scale and spatial position. “Upper brother”, “left brother”, and “upper left brother” refer to a wavelet coefficients which are horizontally, vertically, or diagonally spatially adjacent at the same scale and orientation. “Parent” and “child” refer to wavelet coefficients of the same spatial position and orientation at successive scales. In order to compare our results to those of Huang and Mumford, we used the same wavelet decomposition (Haar) and plotted distributions of the same pairs of coefficients. We obtained contour plots with structure similar to those of Huang and Mumford. Huang and Mumford found that cross sections through the origin of these two-dimensional histograms can be accurately fit with generalized Laplacian curves. While we also found this to be the case for most cross sections, we found some cross sections with shapes reminiscent of the marginal distribution of Figure 5a.

Haar wavelets often perform poorly in practical image processing because of artifacts associated with their lack of smoothness and their lack of localization in the frequency domain. In order to test whether the structure of the distributions of Figure 10 is an artifact of the Haar basis, we repeated the decomposition using an eight-tap quadrature mirror filter [12]. The contour plots maintained their structure. Figure 11 illustrates the two plots that changed the most, corresponding to the second row of Figure 10.

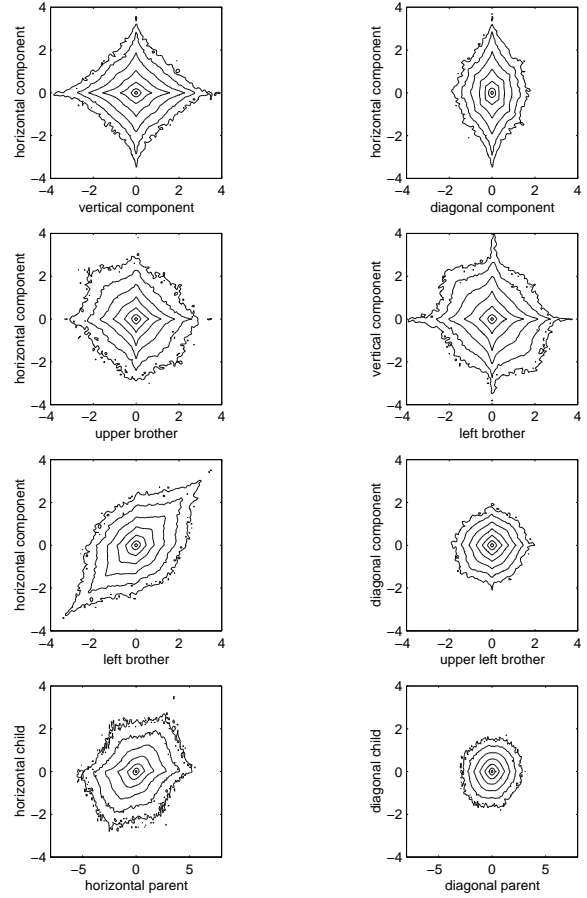


Figure 10: Contour plot of the log(histogram) of finest scale wavelet coefficient pairs for the Teller data set. We have chosen coefficient pairs corresponding to those which Huang and Mumford computed for their set of photographs. Note that the horizontal axes of the bottom two plots are compressed by a factor of two.




Figure 11: Contour plot of the log(histogram) of coefficient pairs corresponding to those in the second row of Figure 10, but using a QMF basis rather than a Haar basis. These are the two contour plots which change most significantly.

9. Discussion

We have found that most of the regularities observed through earlier studies of low dynamic range, restricted view field photographs carry over to natural illumination maps, providing a solid foundation for statistical models of illumination. However, there are some significant differences. Illumination statistics are significantly non-stationary due to their elevation dependence. The presence of bright point sources can significantly alter the power spectrum, so that a $k/f^{2+\eta}$ model does not suffice in general for natural illumination unless it is passed through a compressive nonlinearity. Generalized Laplacian distributions model the distributions of wavelet coefficients and derivatives reasonably well, but the fits are not as close as those observed for more typical photographs.

In order to describe fully the statistics of illumination, one must understand how an illumination maps changes as the camera recording it moves in space. In other words, one must analyze the five-dimensional plenoptic function, which describes all the rays of light passing through every point in a three-dimensional volume [1]. Because image-based rendering involves resampling the plenoptic function [16], statistical priors on this function could enable image-based rendering with sparse data.

References

- [1] E. H. Adelson and J. R. Bergen. The plenoptic function and the elements of early vision. In M. Landy and J. A. Movshon, editors, *Computational Models of Visual Processing*. MIT Press, Cambridge, MA, 1991.
- [2] Frank Canters and Hugo Declair. *The world in perspective: a directory of world map projections*. John Wiley & Sons, New York, 1989.
- [3] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. *Computer Graphics (SIGGRAPH)*, 2000.
- [4] P. E. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. *Computer Graphics (SIGGRAPH)*, 1998.
- [5] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. *Computer Graphics (SIGGRAPH)*, pages 369–78, 1997.
- [6] R. O. Dror, E. H. Adelson, and A. S. Willsky. Estimating surface reflectance properties from images under unknown illumination. In *SPIE Conference on Human Vision and Electronic Imaging*, San Jose, CA, 2001.
- [7] D. Field. Relations between the statistics of natural images and the response properties of cortical cells. *J. Optical Society of America A*, 4:2379–94, 1987.
- [8] D. J. Heeger and J. R. Bergen. Pyramid-based texture analysis/synthesis. *Computer Graphics (SIGGRAPH)*, 1995.
- [9] J. Huang, A. Lee, and D. Mumford. Statistics of range images. In *Proc. Conf. Computer Vision and Pattern Recognition*, Hilton Head, SC, 2000.
- [10] J. Huang and D. Mumford. Statistics of natural images and models. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 541–7, Fort Collins, CO, 1999.
- [11] T. Inui, Y. Tanabe, and Y. Onodera. *Group Theory and Its Applications in Physics*. Springer, Berlin, Heidelberg, second, corrected printing edition, 1996.
- [12] J. D. Johnston. A filter family designed for use in quadrature mirror filter banks. In *Proc. ICASSP*, pages 291–294, 1980.
- [13] S. B. Laughlin. A simple coding procedure enhances a neuron’s information capacity. *Z. Naturforsch.*, 36c:910–912, 1981.
- [14] S. R. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, Ithaca, NY, 1998.
- [15] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. Int. Conf. Computer Vision*, Vancouver BC, 2001.
- [16] L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. *Computer Graphics (SIGGRAPH)*, pages 39–46, 1995.
- [17] J. Portilla and E. P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *IJCV*, pages 49–71, 2000.
- [18] J. Portilla, V. Strela, M. Wainwright, and E. Simoncelli. Adaptive wiener denoising using a gaussian scale mixture model in the wavelet domain. In *ICIP*, 2001.
- [19] D. L. Ruderman. The statistics of natural images. *Network: Comput. Neural Syst.*, 5:517–48, 1994.
- [20] P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. *Computer Graphics (SIGGRAPH)*, 1995.
- [21] E. Simoncelli. Modeling the joint statistics of images in the wavelet domain. In *Proc. SPIE*, volume 3813, 1999.
- [22] E. P. Simoncelli and E. H. Adelson. Noise removal via bayesian wavelet coring. In *International Conference on Image Processing*, 1996.
- [23] S. Teller, M. Antone, M. Bosse, S. Coorg, M. Jethwa, and N. Master. Calibrated, registered images of an extended urban area. *IJCV*, submitted.
- [24] J. H. van Hateren and A. van der Schaaf. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc. R. Soc. Lond. B*, 265:359–366, 1998.