

## Statistics of Real-World Illumination

Ron O. Dror, Thomas K. Leung, Edward H. Adelson, and Alan S. Willsky  
Artificial Intelligence Laboratory and Laboratory for Information and Decision Systems  
Massachusetts Institute of Technology  
rondror@ai.mit.edu, leungt@cs.berkeley.edu, adelson@psyche.mit.edu, willsky@mit.edu


### Abstract

While computer vision systems often assume 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 using 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 natural image statistics literature. However, illumination maps differ from standard photographs in that illumination maps are statistically non-stationary and may contain localized light sources that 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

Computer vision systems have traditionally relied on idealized models of illumination, such as a single point light source or a uniform hemispherical source. Real-world illumination, on the other hand, is highly complex. Surfaces are illuminated not only by luminous sources such as the sun, sky, or indoor lights, but also by light reflected from other surfaces in the environment.

The appearance of a surface depends on its illumination as well as its geometry and reflectance properties. Highly specular surfaces such as chrome or shiny plastic reflect the detailed spatial structure of their illumination. Estimating surface geometry or surface reflectance from images under unknown illumination is therefore difficult. Algorithms for



**Figure 1. (A) A shiny sphere rendered under illumination by a point light source. (B) The same sphere rendered under photographically-acquired real-world illumination. Traditional algorithms for recovery of shape and reflectance, which are designed for the conditions of A, tend to fail on B. Humans perceive reflectance properties more accurately in B ([8]).**

shape from shading or surface reflectance recovery often address this problem by assuming that all light radiates from a single point source. While these algorithms may work properly for images taken in a laboratory, they often fail in the real world (Figure 1). We would like to overcome these limitations by designing computer vision systems based on an accurate statistical description of real-world illumination. This paper takes a first step in that direction, describing statistical properties of real-world illumination based on empirical observation.

In a similar spirit, researchers have recently devoted a great deal of effort to capturing statistics of natural optical images [11, 7, 21, 19]. They 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. These statistics have proven useful for explaining the architecture of biological visual systems [7, 14], developing efficient image coding schemes [21], denoising images [18], and providing models

for probabilistic or Bayesian modeling.

One can measure the illumination 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 map describing illumination at that point. Such spherical images are used as environment maps in computer graphics [3]. If all sources of direct and indirect illumination are relatively distant, the illumination map changes slowly as the hypothetical camera moves through space.

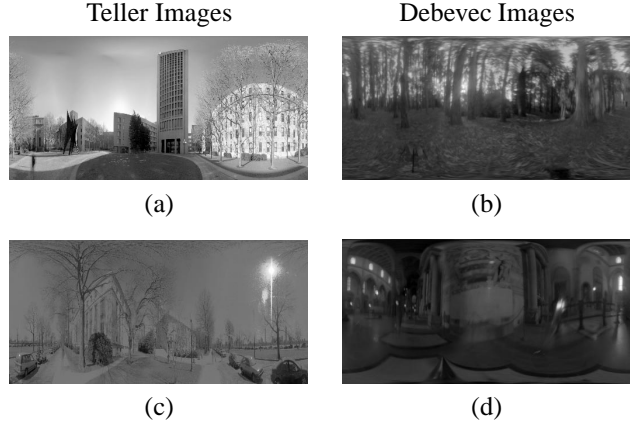
An illumination map is a type of image. However, accurate real-world illumination maps differ from the photographs studied in the natural image statistics literature in several regards. First, illumination maps cover a much wider view angle, spanning the entire sphere instead of a narrow view angle near the horizontal. 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.

We determine the similarities and differences between the statistics of previously studied photographs and those of high dynamic range, photographically-acquired illumination maps. Working with two sets of illumination maps, we 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 [11], 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 encounter 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 that can dominate power spectral measures.

These results might be used as priors to improve the recovery of illumination fields from incomplete data for image-based rendering in computer graphics, as in [15]. An accurate description of illumination statistics may facilitate the development of shape-from-shading algorithms that function robustly under complex, unknown illumination.

The statistics of real-world illumination also play a major role in the recognition of surface reflectance properties from an image, both by humans [8] and by machines [6, 25]. Our work provides a foundation for statistical models of illumination to improve algorithms for reflectance recovery



**Figure 2. 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. Dynamic range has been compressed for display purposes.**

and material recognition.


## 2. Datasets

We worked with two different data sets, each consisting of high dynamic range images that represent the radiance incident 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 9 maps from Debevec’s Light Probe Image Gallery (<http://www.debevec.org/Probes/>) [4]. 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 2.

The images in both data sets were acquired by combining photographs at multiple exposures to obtain pixel values that 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-angle 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




**Figure 3.** 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.

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 3). 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.<sup>1</sup> Figure 2 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 4 shows total illumination intensity distributions for the 95 Teller images and for the 9 Debevec images. Huang and Mumford [11] noted asymmetry in their single pixel distribution due to the presence of sky in many of their images. Our distributions exhibit

<sup>1</sup>In particular, an infinitesimal patch on the sphere at latitude  $\theta$  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 4.** 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.

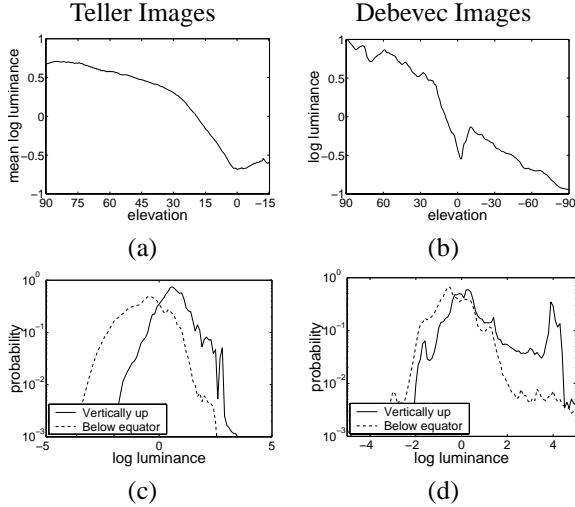
more striking asymmetries, partly because both the Teller and Debevec data sets contain not only sky but other localized light sources. The distribution for the Teller set is particularly asymmetric due to the presence of the sun in many images and to undersaturation in the imaging system at very low 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 power spectra or wavelet statistics (Sections 7 and 8).

Following Huang and Mumford [11], we computed the standard deviation  $\sigma$ , skewness  $S$ , kurtosis  $\kappa$ , and differential entropy  $\mathcal{H}$ .<sup>2</sup> 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  $\mathcal{H} = 2.21$ . 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.


##### 4.1. Non-stationarity

Illumination statistics vary with direction. Figure 5(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 2b). Panels (c) and (d) of Figure 5 each show two histograms at different ranges of elevations. The histograms for higher view directions have a larger mean as well as heavier positive tails, reflecting the larger probability of bright localized

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



**Figure 5. Dependence of illumination on elevation. (a) and (b) show mean luminance as a function of elevation. (c) and (d) each show two histograms of illumination intensities, one for directions within  $30^\circ$  of the upward vertical and the other for directions from  $0^\circ$  to  $15^\circ$  below the equator.**




**Figure 6. Solid lines indicate distribution of horizontal derivatives for log luminance illumination maps. Dashed lines are generalized Laplacian fits.**

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 [10].

## 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



**Figure 7. Joint histogram of horizontally adjacent positions in Teller's illumination maps.**

chosen such that the equator is divided into 512 parts.

Figure 6 shows the resulting marginal distributions of horizontal derivatives for the two datasets. Both distributions, like the distribution reported by 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  $\kappa = 17.4$ . As 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. 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  $\mathcal{P}(x) \propto \exp(-|x/s|^\alpha)$  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 6). 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. This difference, as well as the difference between the distributions of our two data sets, may stem from the different distributions of luminous sources.

## 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 that are horizontally adjacent as defined in Section 5. Figure 7 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 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 4.

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 0.103 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

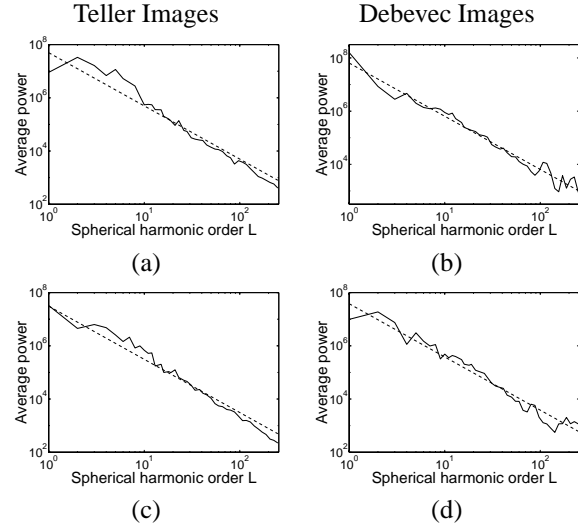
Much early work on natural image statistics focused on the regularity of power spectra. A number of authors [7, 19] have observed that two-dimensional power spectra of natural images typically fall off as  $1/f^{2+\eta}$ , where  $f$  represents the modulus of the frequency and  $\eta$  is a small constant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

The natural spherical equivalent of the planar Fourier transform 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 [12].

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 [12]. 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 8 shows the relationship between average power and harmonic order for the four illumination maps of Figure 2, when pixel value is proportional to log luminance. All four images have power spectra that 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.



**Figure 8. Spherical harmonic power spectra (solid lines) of illumination maps (a), (b), (c), and (d) in Figure 2, with pixel value proportional to log luminance. The dotted lines of slope  $-2$  correspond to power spectra of the form  $k/L^2$ .**

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 2a and b, which lack concentrated primary light sources, have spherical harmonic spectra that are well approximated by  $k/L^{2+\eta}$  with  $\eta$  small. On the other hand, illumination maps that contain intense, localized light sources have smooth power spectra that remain flat at low frequencies before falling off sharply at higher frequencies. The illuminations of Figure 2c and d both display this behavior; the power spectrum of a linear luminance version of Figure 2c is shown in Figure 9. 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 9).

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, on the other hand, may be dominated by light sources occupying a small spatial area. Once the relative strength of such sources is



**Figure 9.** 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. Clipping these extremely bright pixels reduces power at all frequencies and produces a more linear power spectrum. The dotted lines of slope  $-2$  correspond to power spectra of the form  $k/L^2$ .


reduced through clipping or a logarithmic transformation, illumination maps have power spectra similar to those of typical photographs.

## 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 [9, 17], or reflectance classification [6].

Previous 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. The plots shown in this section are based on equal area cylindrical projections of spherical log-luminance illumination maps, with  $k = \frac{2}{\pi}$ . We experimented with several projections of the sphere to the plane for the purpose of wavelet analysis and obtained similar results.

Figure 10a 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

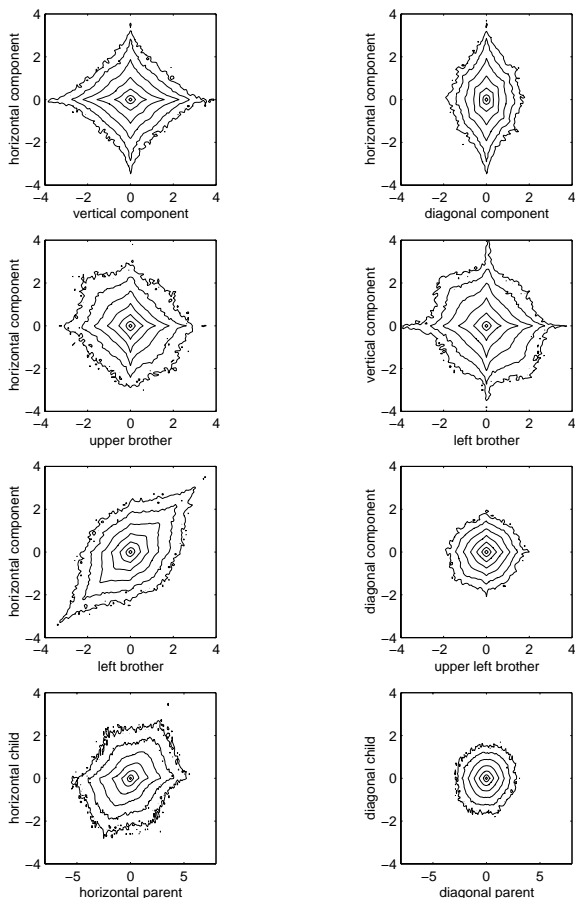


**Figure 10.** Distributions of Haar wavelet coefficients at successive scales (thick lines), along with generalized Laplacian fits (thin lines), for the 95 Teller images. Left, horizontal bands. Right, vertical bands. Parameters of fits are: horizontal, finest scale,  $\alpha = 0.58$ ,  $s = 0.043$ ; second-finest scale,  $\alpha = 0.52$ ,  $s = 0.056$ ; third-finest scale,  $\alpha = 0.50$ ,  $s = 0.10$ ; vertical, finest scale,  $\alpha = 0.54$ ,  $s = 0.043$ ; second-finest scale,  $\alpha = 0.48$ ,  $s = 0.052$ ; third-finest scale,  $\alpha = 0.43$ ,  $s = 0.067$ .


and fits the generalized Laplacian closely. The marginal distributions increase in variance at successively coarser scales. Figure 10b shows the corresponding distributions for vertically oriented wavelets. The generalized Laplacian fit is poorest at the finest scale; this particular distribution corresponds approximately to the distribution of horizontal pixel-wise differences in Figure 6a. The poor fit may be due to noise at the finest scale.

Figure 11 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 wavelet coefficients that 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 theirs. 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 6a.

Haar wavelets often perform poorly in practical image processing because of artifacts associated with their lack of smoothness and lack of localization in the frequency domain. To test whether the structure of the distributions of



**Figure 11. 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 that 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 12. Contour plot of the log(histogram) of coefficient pairs corresponding to those in the second row of Figure 11, but using a QMF basis rather than a Haar basis. These are the two contour plots that depend most significantly on the basis.**

Figure 11 is an artifact of the Haar basis, we repeated the decomposition using an eight-tap quadrature mirror filter [13]. The contour plots maintained their structure. Figure 12 illustrates the two plots that changed the most.

## 9. Discussion

We have found that most of the regularities observed through earlier studies of low dynamic range, restricted field-of-view photographs carry over to real-world illumination maps, providing a solid foundation for statistical models of illumination. These similarities are accompanied by several 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 the illumination map is passed through a compressive non-linearity. 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.

These conclusions have important implications for robust estimation of shape, illumination, and reflectance from images. For example, in [6] we consider the problem of differentiating between surfaces of different reflectance, such as chrome, shiny plastic, and paper, on the basis of a single image under unknown illumination. The fact that real-world illumination possesses greater regularity in the wavelet domain than in the frequency domain helps guide the choice of image features for use in classifying images according to surface reflectance.

We hope to extend our statistical description of illumination by considering the statistics of the five-dimensional plenoptic function that describes all the rays of light passing through every point in a three-dimensional volume [1]. In addition to understanding the statistics of illumination

maps, we wish to understand how an illumination map changes as the camera recording it moves in space. Because image-based rendering involves resampling the plenoptic function [16], statistical priors on this function could facilitate image-based rendering with sparse data.

## Acknowledgments

Seth Teller, Neel Master, and Michael Bosse shared the data set from the MIT City Scanning Project and helped us use it to construct illumination maps. This work was supported by NDSEG and Whitaker Fellowships to R.O.D., by NIH Grant EY11005-04 to E.H.A., by a grant from NTT to the MIT Artificial Intelligence Lab, by a contract with Unilever Research, and by ONR Grant N00014-00-1-0089 to A.S.W.

## 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] F. Canters and H. Declair. *The World in Perspective: A Directory of World Map Projections*. John Wiley & Sons, New York, 1989.
- [3] 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.
- [4] P. E. 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.
- [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. Surface reflectance estimation and natural illumination statistics. In *Proc. of IEEE Workshop on Statistical and Computational Theories of Vision*, Vancouver, Canada, July 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] R. Fleming, R. O. Dror, and E. H. Adelson. How do humans determine reflectance properties under unknown illumination? In *Proceedings of CVPR Workshop on Identifying Objects Across Variations in Lighting: Psychophysics and Computation*, 2001.
- [9] D. J. Heeger and J. R. Bergen. Pyramid-based texture analysis/synthesis. *Computer Graphics (SIGGRAPH)*, 1995.
- [10] J. Huang, A. Lee, and D. Mumford. Statistics of range images. In *Proc. Conf. Computer Vision and Pattern Recognition*, Hilton Head, SC, 2000.
- [11] 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.
- [12] T. Inui, Y. Tanabe, and Y. Onodera. *Group Theory and Its Applications in Physics*. Springer, Berlin, Heidelberg, second, corrected printing edition, 1996.
- [13] J. D. Johnston. A filter family designed for use in quadrature mirror filter banks. In *Proc. Int. Conf. Acoustics, Speech, and Signal Processing*, pages 291–294, 1980.
- [14] S. B. Laughlin. A simple coding procedure enhances a neuron’s information capacity. *Z. Naturforsch.*, 36c:910–912, 1981.
- [15] S. R. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, Ithaca, NY, 1998.
- [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. *Int. J. Computer Vision*, 40:49–71, 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 *Proc. Int. Conf. Image Processing*, 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 *Proc. Int. Conf. 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. *Int. J. Computer Vision*, 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.
- [25] Y. Weiss. Deriving intrinsic images from image sequences. In *Proc. Int. Conf. Computer Vision*, 2001.