Statistics of Real-World Illumination

Ron O. Dror & Edward H. Adelson

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: To describe the spatial structure of illumination in real-world environments.

Motivation: While research in computer vision often assumes simple illumination models, real-world illumination is highly complex. Illumination includes not not only direct illumination from the sun, sky, or indoor lights, but also indirect illumination by light reflected from other surfaces in the environment.

The appearance of a surface depends not only on its geometry and reflectance properties, but also on its illumination. Estimating surface geometry or surface reflectance from a single image is difficult when illumination is unknown. Traditional computer vision systems have addressed this problem by assuming simple forms of illumination such as point source illumination. Such systems do not operate robustly in the real world because their assumptions are often false. We would like to overcome these limitations by designing computer vision systems based on an accurate statistical description of real-world illumination.

Previous Work: Researchers have devoted a great deal of effort to capturing the statistics of "natural images," typical photographs of indoor or outdoor scenes. These studies have found that such photographs display a great deal of regularity, particularly in power spectra and distributions of wavelet coefficients [3, 5, 4].

Approach: One can measure the illumination incident from every direction at a point in the real world using a camera located at the point of interest. By combining photographs from that camera in every direction, one can compose a spherical illumination map describing illumination at that point. Figure 1 shows two examples.

Figure 1: Examples of illumination maps. These spherical maps are shown in equal-area cylindrical projection. The outdoor image at left is drawn from a data set due to Teller [6], while the indoor image at right comes from a set due to Debevec [1]. Pixel values have been passed through a compressive nonlinearity for display purposes.

An illumination map is a type of image. We therefore expect that illumination maps may share some of the statistical regularities reported in the literature on natural image statistics. We analyze marginal and joint distributions of illumination intensity, distributions of directional intensity derivatives, spherical harmonic power spectra, and marginal and joint wavelet coefficient distributions for a set of illumination maps. For example, Figure 2 shows probability distributions of wavelet coefficients at several scales and orientations. Although the statistics of illumination maps share a number of similarities with those of typical photographs, we have discovered several significant differences.

Difficulty: While previously studied photographs cover a narrow view angle near the horizontal, illumination maps ideally cover the entire sphere. This complicates spectral analysis and wavelet analysis of the illumination maps. For example, we use the spherical harmonic decomposition in place of the traditional Fourier transform.

Accurate illumination maps also possess a much higher dynamic range than previously analyzed photographs.

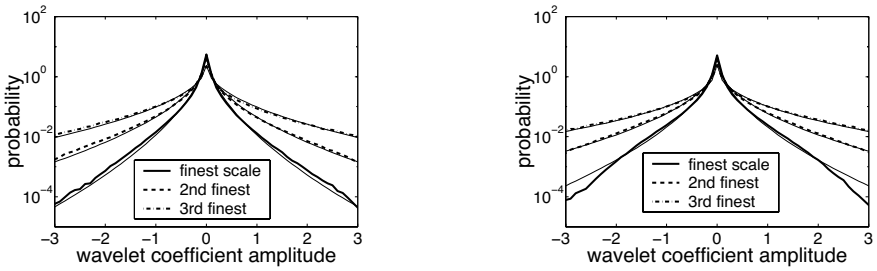

Figure 2: Histograms of Haar wavelet coefficients at successive scales (thick lines), along with generalized Laplacian fits (thin lines). Left, horizontal wavelet bands. Right, vertical wavelet bands. The distributions are similar for each of the illumination maps analyzed.

In fact, they may contain localized primary light sources such as incandescent lights or the sun itself. We therefore work with illumination maps acquired by compositing a series of photographs taken at different exposures.

Impact: An understanding of the regularities in real-world illumination may lead to more practical and robust computer vision systems. We have applied our initial results to the recognition of surface reflectance properties from images under unknown lighting [2], an important step in allowing computer vision systems to recognize materials such as metal, plastic, or paper. Shape-from-shading algorithms based on an accurate description of illumination statistics could function robustly under complex, unknown illumination. A statistical description of illumination could also improve the recovery of illumination fields for image-based rendering in computer graphics.

Future Work: So far, we have studied the statistical properties of illumination at individual points in space. As one moves across a surface, the illumination map changes; for example, one might move from a shadow into bright sunlight. By analyzing multiple illumination maps acquired at known distances from one another, we hope to describe the nature of these variations and to produce a more complete statistical model of illumination.

Research Support: This research is supported by NTT, Unilever, an NIH Grant EY11005-04 to E.H.A., and an NDSEG Fellowship to R.O.D.

References:

[1] 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.

[2] 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.

[3] D. J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A*, 4:2379–94, 1987.

[4] J. Huang and D. Mumford. Statistics of natural images and models. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 541–47, 1999.

[5] D. L. Ruderman. The statistics of natural images. *Network: Comput. Neural Syst.*, 5:517–48, 1994.

[6] S. Teller, M. Antone, M. Bosse, S. Coorg, M. Jethwa, and N. Master. Calibrated, registered images of an extended urban area. *International Journal of Computer Vision*, submitted.