Estimating Surface Reflectance Properties from Images under Unknown Illumination

Ron O. Dror[a], Edward H. Adelson[b], and Alan S. Willsky[a]

[a]Department of Electrical Engineering and Computer Science
[b]Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology

ABSTRACT

Physical surfaces such as metal, plastic, and paper possess different optical qualities that lead to different characteristics in images. We have found that humans can effectively estimate certain surface reflectance properties from a single image without knowledge of illumination. We develop a machine vision system to perform similar reflectance estimation tasks automatically. The problem of estimating reflectance from single images under unknown, complex illumination proves highly underconstrained due to the variety of potential reflectances and illuminations. Our solution relies on statistical regularities in the spatial structure of real-world illumination. These regularities translate into predictable relationships between surface reflectance and certain statistical features of the image. We determine these relationships using machine learning techniques. Our algorithms do not depend on color or polarization; they apply even to monochromatic imagery. An ability to estimate reflectance under uncontrolled illumination will further efforts to recognize materials and surface properties, to capture computer graphics models from photographs, and to generalize classical motion and stereo algorithms such that they can handle non-Lambertian surfaces.

Keywords: reflectance, illumination, estimation, BRDF, natural image statistics

1. INTRODUCTION

A human observer differentiates between a mirrored surface and a white matte surface effortlessly at a single glance. While most people take this ability for granted, the task is nontrivial from a computational point of view. Figure 1 shows two images of each of four different spheres. Because each sphere was photographed under two different illumination conditions, pixels in one image of a given sphere are quite different from corresponding pixels in the other image of the same sphere. In principle, all four spheres could be perfect chrome reflectors. Because a mirror simply reflects its illumination environment, a properly illuminated chrome sphere could produce an image identical to that of a white matte sphere.

Despite this ambiguity, humans easily recognize that the images in each column represent spheres of similar materials, while the images in different columns represent objects of different materials. In fact, recent psychophysical work has shown that humans are capable of much more subtle distinctions in reflectance properties on the basis of images of isolated surfaces which provide no independent information about illumination.[1] Presumably the human visual system relies on prior information in order to conclude that, although each image could be achieved through several combinations of surface reflectance, environmental illumination, and surface geometry, some explanations are much more likely than others. This paper establishes a statistical framework for reflectance estimation under unknown, real-world illumination.

Several practical applications in computer vision and graphics motivate our efforts. First, an ability to estimate reflectance under unknown illumination facilitates visual material recognition, because different physical surfaces such as metal, plastic, and paper possess different optical reflectance properties. Second, reconstruction of a scene from photographs for computer graphics requires inference of both the geometry and the reflectance of visible surfaces. Third, an ability to estimate reflectance from image data under unknown lighting conditions may help overcome the

Further author information: (Send correspondence to R.O.D.)
R.O.D.: rondror@ai.mit.edu; MIT 35-427, 77 Massachusetts Ave., Cambridge, MA 02139
E.H.A.: adelson@psyche.mit.edu; MIT NE20-444, 77 Massachusetts Ave., Cambridge, MA 02139
A.S.W.: willsky@mit.edu; MIT 35-433, 77 Massachusetts Ave., Cambridge, MA 02139

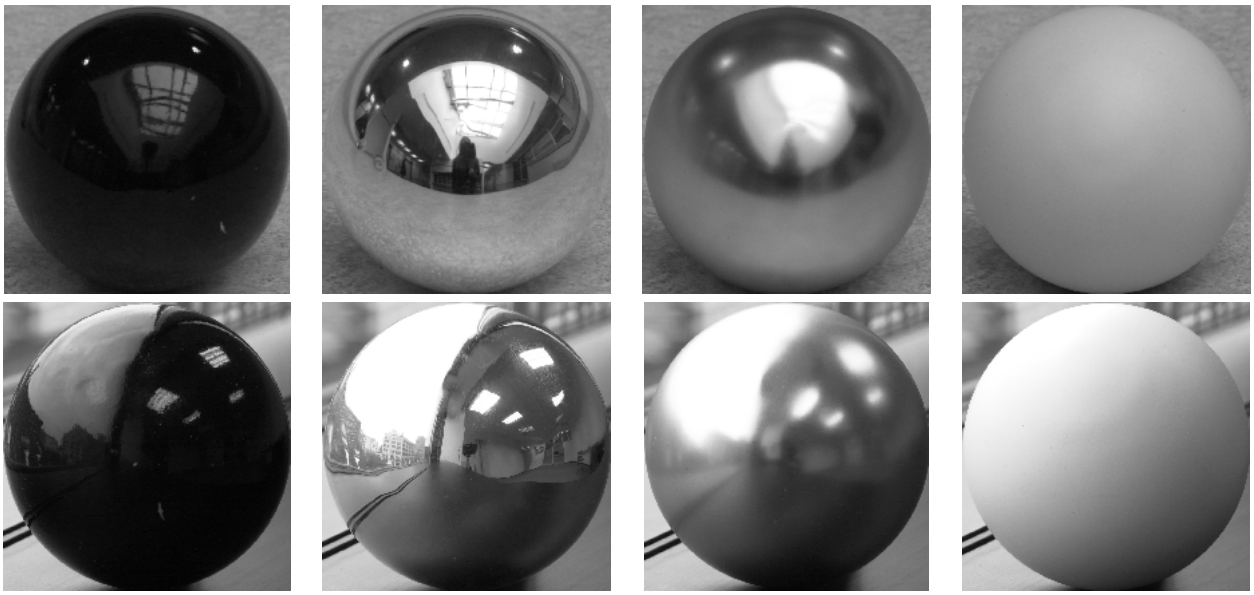

Figure 1. The two rows show photographs of the same four spheres under two different illuminations. Each sphere possesses distinct reflectance properties which a human can recognize under different illuminations: (A) shiny black plastic, (B) chrome, (C) rough metal, (D) white matte paint.

limitations of shape-from-shading algorithms which assume that reflectance is known in advance, and of classical algorithms for motion or stereo estimation, which assume Lambertian surface reflectance.

The importance of reflectance models in computer graphics has motivated several researchers to develop image-based reflectance estimation techniques. Tominaga et al.[2] present a method for estimating Phong model parameters from an image of a uniform cylindrical surface. They require illumination by a point light source, although they estimate the exact location of the light source from image data. Sato et al.[3] as well as Marschner[4] develop similar techniques which accommodate a more general geometry acquired through laser range scanning. These methods, unlike ours, do not apply to photographs taken in the natural world under complex lighting conditions.

The inverse global illumination techniques of Yu et al.[5] and Yu and Malik[6] do handle reflectance estimation without fully specified illumination, but with input data requirements much greater than those of the methods in this paper. Given a collection of photographs of a local scene (and, when appropriate, a light probe image representing illumination from distant sources), the inverse global illumination algorithm iteratively estimates both the illumination and reflectance of every surface patch in the local scene. This requires a collection of photographs representing all surfaces which cast light upon the surfaces of interest, and in which all the primary light sources are identifiable. Our approach, on the other hand, requires only an image of the surface whose reflectance is in question. We avoid estimating illumination explicitly by characterizing it statistically.

We estimate reflectance from single monochrome images of the surface of interest, without using contextual cues from the surrounding image. In practice, a computer vision system may combine a number of cues to estimate reflectance properties. Color spectral techniques can separate specular and diffuse reflectance components for dielectrics.[7,2] Motion cues differentiate between diffuse and specular reflectance. Visual context provides partial information about illumination conditions. We have found, however, that humans can effectively estimate certain surface reflectance properties even in the absence of these cues.[1] We chose to work with single, monochrome images, because we wish to determine what information the basic image structure captures about reflectance. Our techniques could be improved by drawing on additional sources of information. We simplify the problem at this stage by assuming that the surface under observation has homogenous reflectance properties and that its geometry is known.

We begin by defining the problem of reflectance estimation under unknown illumination mathematically and showing that this problem is ill posed (Section 2). We then examine available prior information about real-world reflectance and illumination (Section 3). In particular, the spatial structure of real-world illumination possesses

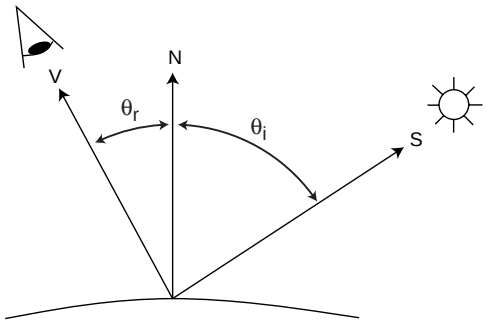

Figure 2. A surface patch with normal N. The BRDF is a function of light source direction S and view direction V.

strong statistical regularities akin to those studied in the natural image literature. We develop a method for using this information to classify images according to surface reflectance (Section 4). Finally, we demonstrate reflectance estimation results on both real and synthetic images (Section 5).

2. MATHEMATICAL PROBLEM STATEMENT

The reflectance of an opaque surface patch can be described as a bidirectional reflectance distribution function (BRDF), which specifies what proportion of the light incident from each possible illumination direction is reflected in each possible view direction. Figure 2 shows a surface patch with normal N illuminated by a light source in direction S and observed by a viewer in direction V. In a three-dimensional world, two angles are necessary to uniquely specify S and two more to specify V. Hence the BRDF is a function of four continuous angular variables, specifying the ratio of the reflected radiance in a particular direction to incident irradiance from a differential solid angle centered on the incident direction. We denote it by $f(\theta_i, \phi_i; \theta_r, \phi_r)$, where θ_i and θ_r are the angles of S and V, respectively, from the surface normal N, and ϕ_i and ϕ_r are their respective azimuthal angles. We wish to estimate f from an observed image.

Given the BRDF of a surface patch and the illumination incident on it from every direction in a hemisphere surrounding its normal, we can compute its apparent brightness to a viewer. If $I(\theta_i, \phi_i)$ gives the radiance of illumination incident from direction (θ_i, ϕ_i), the total reflected radiance of the surface patch in the view direction (θ_r, ϕ_r) is given by[8]

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

This expression highlights the ill-posed nature of the reflectance estimation problem. We wish to estimate the BRDF, a continuous function of four variables, from a single image of the surface. Each pixel of the image depends not only on the BRDF, but on the unknown illumination, a function of two variables which may differ from one point on the surface to another. The observed image is a function of only two variables. Even assuming that geometry is known and that the BRDF is the same at all points on the surface, the unknowns occupy a much larger space than the observations. In order to choose the most likely reflectance given the observed image, we must exploit prior information on real-world reflectance and illumination.

3. STATISTICAL PRIOR INFORMATION

3.1. Reflectance

Basic physics requires that the BRDF of a passive material be normalized, meaning that for any irradiation, total reflected energy is less than or equal to total incident energy. A physical BRDF must also satisfy the reciprocity principle, which guarantees symmetry between incident and reflected directions ($f(\theta_i, \phi_i; \theta_r, \phi_r) = f(\theta_r, \phi_r; \theta_i, \phi_i)$). The space of physically realizable BRDFs remains huge, including rare reflectances such as those of holograms.

In the real world, some reflectances are far more common than others. An ideal reflectance estimator would exploit a probability distribution over the entire space of realizable BRDFs. Unfortunately, we lack the data to

formulate such a distribution. Instead, we rely on one of two cheaper alternatives. The examples of Section 5 assume that, given a finite but arbitrary set of candidate reflectance functions, we wish to identify which one most closely represents an observed surface. This simplification reduces the infinite-dimensional estimation problem to a discrete classification problem.

Alternatively, one might restrict the reflectances under consideration to those characterized by a parametric model from computer graphics. One must then estimate the parameter values that most closely match the reflectance properties of the observed surface. This approach reduces the estimation problem to a regression problem in a finite number of variables. Although no existing parametric model captures all real-world reflectances accurately, computer graphics researchers have painstakingly designed models to produce reasonable rendering results for most materials. We have worked primarily with Ward's isotropic reflectance model,[9,10] which takes the form

$$f(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{\rho_d}{\pi} + \rho_s \frac{1}{\sqrt{\cos\theta_i \cos\theta_r}} \frac{\exp[-\tan^2 \delta/\alpha^2]}{4\pi\alpha^2}, \tag{2}$$

where δ is the angle between the surface normal and a vector bisecting the incident and reflected directions. The free parameters of this model are ρ_d, the fraction of incident energy reflected by the diffuse (Lambertian) component, ρ_s, the fraction of energy reflected by the specular component, and α, surface roughness measured as the standard deviation of surface slope. Higher α implies a more blurred specular component.

3.2. Illumination

In principle, illumination could vary almost arbitrarily from one point on a surface to another. In practice, however, the illuminations of two nearby surface points are usually similar. In fact, if all sources of direct and indirect illumination are sufficiently distant and the surface is convex, one can represent surface illumination by a single spherical map defining the intensity of light incident from each real-world direction. The illumination of each surface patch is then given by a hemisphere of this global map. One can conveniently capture a spherical illumination map for any point in the real world photographically using a fisheye lens or a mirrored ball. Debevec's Light Probe Image Gallery[11] consists of several illumination maps acquired with high dynamic range from a combination of photographs at multiple exposures. Under the distant illumination assumption, one can use such an illumination map to render an object as it would appear at that location.[12]

The brightness of a point on a convex surface under distant illumination depends only on the local surface orientation. Hence an image of the surface determines a relationship between surface orientation and observed radiance for a particular illumination condition. One can therefore apply similar reflectance estimation techniques to any convex object whose geometry is known or measured in advance. We have chosen to illustrate our techniques using spheres because we can capture their geometry precisely from image contours.

Photographically acquired real-world illumination maps possess a great deal of predictable statistical structure in the frequency and wavelet domains. This observation about illumination parallels a growing body of research which has demonstrating that "natural" images in both outdoor and indoor environments possess predictable statistical structure. For example, two-dimensional power spectra tend to fall off as an inverse power of frequency, varying little from image to image.[14] Wavelet domain analysis has proven particularly powerful in capturing natural image structure. Distributions of wavelet coefficients at any given scale and orientation are heavy-tailed, falling off much more slowly than a Gaussian distribution.[15] The variance of wavelet coefficient distributions tends to increase in a geometric sequence as one moves to successively coarser scales. Joint distributions of wavelet coefficients across scales capture further statistical structure.[16] Figure 3 illustrates some of these wavelet domain characteristics for two illumination maps.

Figure 4 shows synthetic images of two identical spheres under different illuminations. Surface reflectance is more easily identified in image B, rendered under a photographically acquired illumination map, than in image A, rendered under point source illumination. Comparison of photographs of a sphere in a normally illuminated room and in a black room with a point light source reveals the same effect. The simplicity of point source illumination violates the statistical regularities of typical natural images. Most previous work in reflectance estimation has considered the case of point source illumination as a convenient starting point. We wish instead to take advantage of the statistical complexity of natural illumination in estimating reflectance.

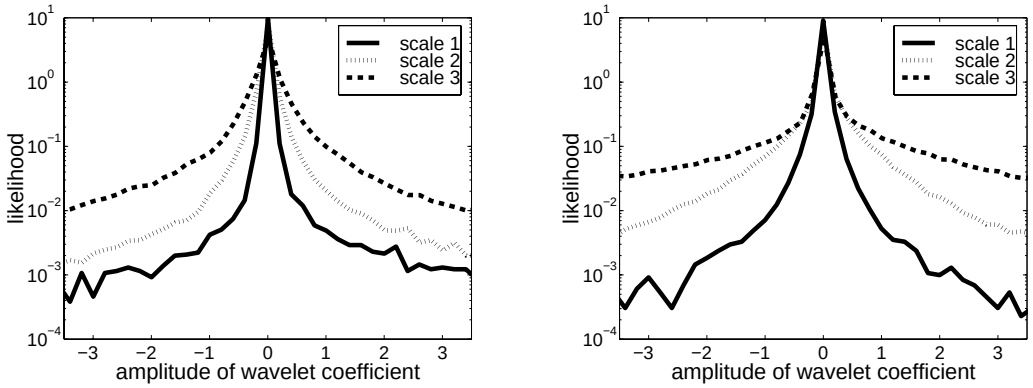

Figure 3. Distributions of wavelet coefficients at one orientation for three successive scales, with scale 1 being the finest, for (A) an indoor illumination map (Galileo's tomb) and (B) an outdoor illumination map (eucalyptus grove), both from Debevec.[11] All the distributions have a characteristic heavy-tailed form, with variance increasing at coarser scales. These distributions were computed as histograms of bands of a nine-tap quadrature mirror filter wavelet pyramid[13] constructed from an unwrapped annulus of Debevec's high dynamic range images (see Section 4.2).

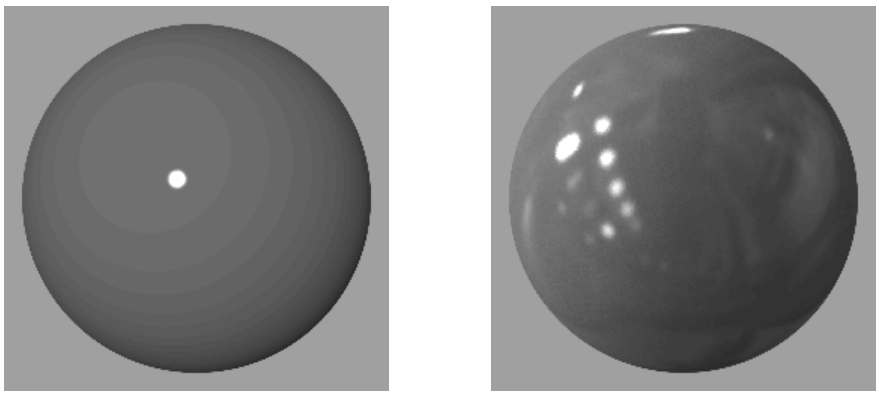

Figure 4. (A) A sphere rendered under illumination by a point light source, with reflectance specified by Ward model parameters $\rho_d = 0.083$, $\rho_s = 0.097$, $\alpha = 0.03$. (B) The same sphere rendered under photographically acquired illumination.

4. A METHOD FOR REFLECTANCE CLASSIFICATION

4.1. Bayesian formulation

The ideal Bayesian approach to reflectance estimation would involve marginalizing over all possible illuminations to find the most likely reflectance for a given observed image:

$$\hat{\nu} = \arg\max_{\nu} P(\nu|R) = \arg\max_{\nu} P(\nu) \int_I P(I)P(R|\nu, I)dI, \tag{3}$$

where ν denotes the parameters of a reflectance function, I denotes illumination as a function of direction, and R denotes the observed image radiance as a function of surface orientation.

Unfortunately, even if one could formulate the prior probability of any given illumination map $P(I)$ explicitly, integration over all possible illuminations is computationally intractable. Replacing the integration with a maximum over illuminations will lead to incorrect reflectance estimates. Consider a photograph of a white matte sphere, corrupted by slight high-frequency imaging noise. One could explain this image approximately as a white matte sphere under any of a number of illuminations, but none of these would predict the noisy image exactly. On the other hand, one could explain the photograph precisely as a chrome sphere under just the right illumination. Thus the single most likely combination of reflectance and illumination might indeed involve a chrome sphere. Integrating over all possible illuminations, however, would reveal that a more likely reflectance is white matte, because for that reflectance a large number of illuminations produce approximately the observed image. Joint estimation of illumination and reflectance, even when feasible, may fail to identify the most likely reflectance. This argument parallels Bayesian approaches such as Freeman's generic viewpoint assumption.[17]

4.2. A machine learning approach

Practical reflectance estimation requires a method to leverage the statistical regularity of real-world illumination while avoiding integration over all possible illuminations. Our approach applies machine learning techniques to determine relationships between image features and reflectance. Intuitively, the feature set should capture notions such as the presence or absence of edges, which we associate with the presence or absence of a sharp specular component because natural illumination maps generally contain sharp edges. Inspired by the observations of Section 3.2 on real-world illumination statistics and by work in texture modeling,[18,19] we chose features based on statistics of the distributions of wavelet coefficients at different scales and orientations, as well as the distribution of pixel intensities in the original image. Heeger and Bergen[18] based their texture representation on distributions of wavelet coefficients and pixel intensities.

Previous work in texture analysis and natural image statistics has assumed stationary image statistics. An image of a sphere will possess nonstationary image statistics because a perspective or orthographic projection will compress features near the edges. One could reduce these effects by considering radiance to be a function of orientation and performing the analysis in a spherical domain*, using spherical wavelet transforms.[20] We have chosen initially to simply warp the observed image into one with more nearly stationary statistics. In particular, our algorithm extracts an annulus of the spherical image and unwraps it into a rectangle using a polar-to-rectangular coordinate transformation (Figure 5).

A two-dimensional pyramid is constructed from this warped image using nine-tap symmetric quadrature mirror filters.[13] We then compute histograms to approximate the distributions of pixel intensities in the original image and the distribution of wavelet pyramid coefficients at each scale and orientation. Each distribution is characterized by a set of numerical statistics including mean, variance, skew, kurtosis, and the 10th, 50th, and 90th percentiles. Figure 5 summarizes feature computation as a flowchart.

Finally, we determine experimentally which of the computed statistics are meaningful and use them to build a classifier. We chose support vector machines (SVMs) for classification because they tend to generalize well given a limited number of training samples and a large number of features.[21] Our implementation utilizes the SVMTorch

*Even when the image is defined as a function on the sphere instead of the plane, its statistics are not stationary. Consider a typical reflectance function consisting of a diffuse and a specular component. A localized light source contributes most to the diffuse component for a surface patch whose normal is in the light source direction. The same light source contributes most to the specular component for a surface patch whose normal bisects the light direction and the view direction. Even if illumination is isotropic, image statistics will vary spatially.

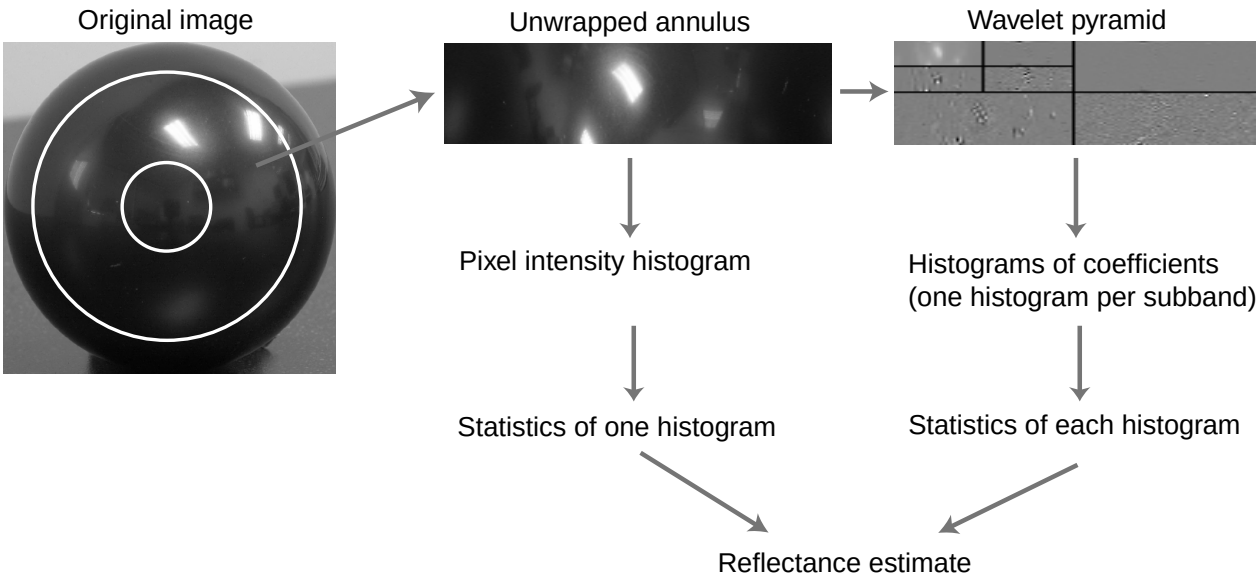

Figure 5. Flowchart for computation of image features, which applies to both testing and training of the classifier. The features are histogram statistics, computed on the original image and on its wavelet transform.

software[22] with Gaussian kernels to train and apply SVM classifiers. Although this paper only presents classification results, the SVM approach generalizes naturally to regression on reflectance parameters. Section 5 further discusses our choice of specific features.

Because our analysis techniques rely solely on the image of the surface of interest, they suffer from ambiguity between the overall strength of illumination and the overall lightness of the surface. A white matte sphere under dim illumination and a gray matte sphere under bright illumination will produce identical images. Resolution of this ambiguity requires contextual information from the remainder of the image or scene. Because color constancy and lightness estimation have been studied separately,[23,24] we eliminate this problem from the current study by normalizing our images for overall strength of illumination, as measured by the brightness of a standard white surface positioned perpendicular to the viewer at the position of the surface under observation.

We train a classifier using either photographs, or synthetic images rendered under photographically acquired illumination maps. Synthetic images have the advantage that the surface BRDF is known exactly. To create synthetic images, we used Ward's *Radiance* package,[10] which efficiently implements the Ward reflectance model. Our rendering methodology parallels that of Debevec.[12] When working with synthetic images, our machine learning process treats the rendering machinery as a black box, using only the final rendered images in the training process.

5. RESULTS

5.1. Synthetic Images

We trained classifiers on synthetic images of spheres of six known reflectance properties. The six reflectances were specified by Ward model parameter settings chosen to correspond to common materials of distinctly different appearances. Each sphere was rendered under nine illuminations specified by Debevec's high dynamic range light probes.[11] These illumination maps represent diverse lighting conditions from four indoor settings and five outdoor settings. Each sphere was rendered from a view angle midway between a top view and a horizontal view. The resulting renderings were converted from *Radiance*'s native high dynamic range format to floating point images for further processing. Figure 6 shows a sample rendering of each sphere, together with the Ward model parameters for each.

To test the accuracy of classification based on different combinations of statistics using only a small total number of data points, we performed a variant of leave-one-out cross-validation. We classified the six images corresponding to each illumination using a classifier trained on the images corresponding to the other eight illuminations. By repeating this process for each of the 9 illuminations, we obtained a total of 54 test cases, one for each rendered image.

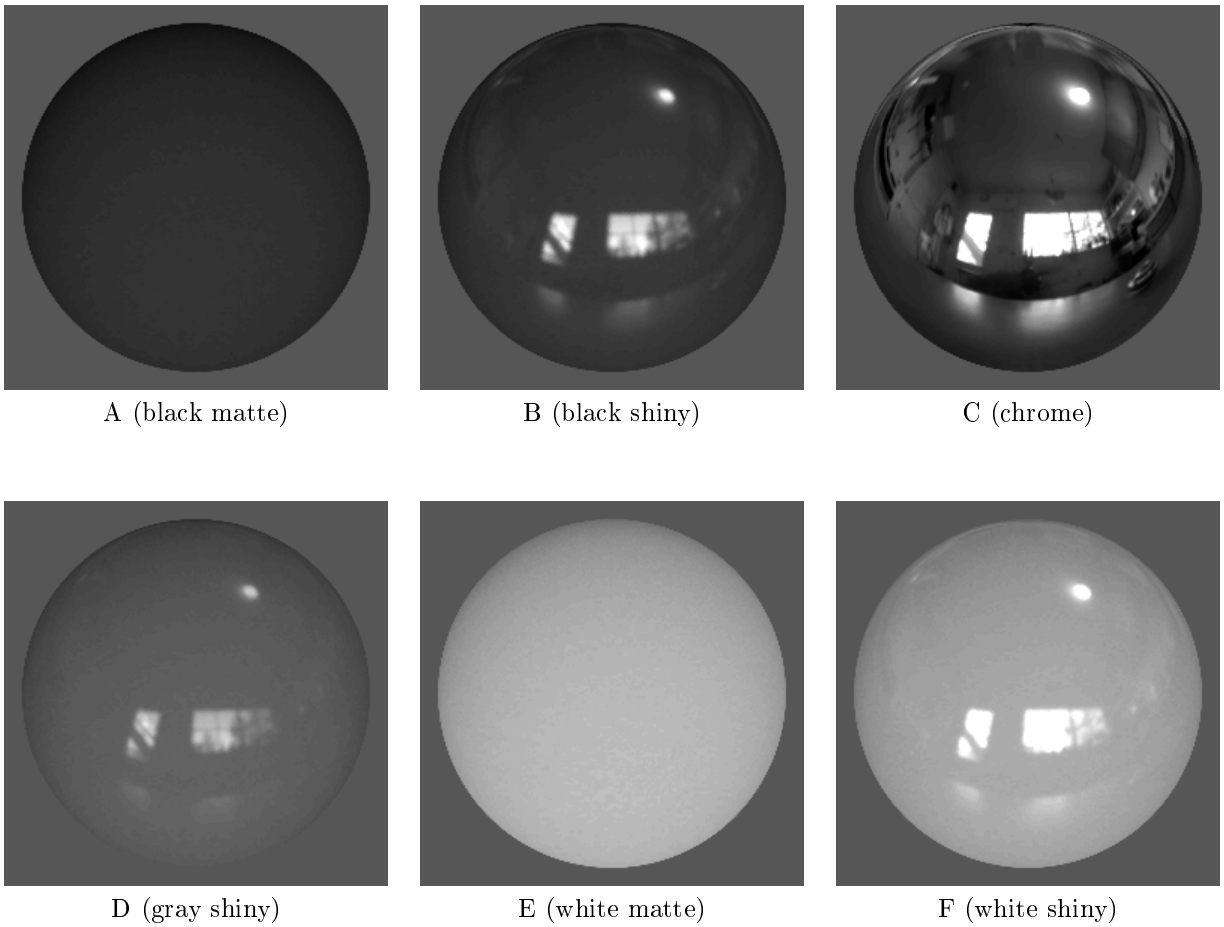

A (black matte) B (black shiny) C (chrome)

D (gray shiny) E (white matte) F (white shiny)

Figure 6. Rendered spheres used in classification experiments. All are shown under the same illumination map. Ward model parameters are as follows: (A) $\rho_d = .1$, $\rho_s = 0$, (B) $\rho_d = .1$, $\rho_s = .1$, $\alpha = .01$, (C) $\rho_d = 0$, $\rho_s = .75$, $\alpha = 0$, (D), $\rho_d = .25$, $\rho_s = .05$, $\alpha = .01$, (E) $\rho_d = .9$, $\rho_s = 0$, (F) $\rho_d = .7$, $\rho_s = .25$, $\alpha = .01$.

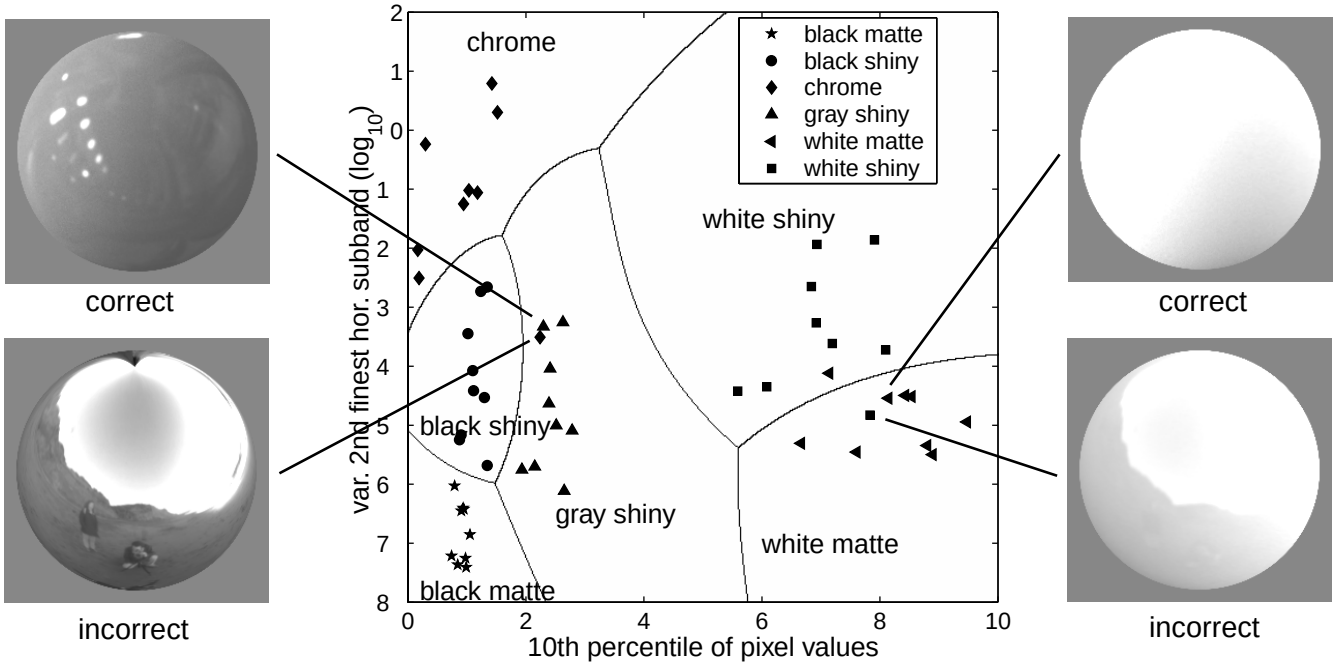

Figure 7. A classifier trained on two features. Solid symbols indicate locations of training samples, as specified by the legend. Lines separate the regions which the classifier assigns to different reflectances. The pair of images on the left-hand side of the figure show a chrome sphere incorrectly classified as gray shiny, and the correctly classified gray shiny sphere falling closest to it in the space defined by these two statistics. Similar, the image pair on the right-hand side include a white shiny sphere incorrectly classified as white matte, and a white matte sphere which falls close to it. The vertical axis measures the variance of wavelet coefficients of the warped image at the second-finest horizontal scale.

Classification based on five or more statistics produced accuracies as high as 53 of 54 correct. For the purposes of visualization, Figure 7 illustrates the behavior of a classifier based on only two image statistics. The horizontal axis denotes the 10th percentile of the distribution of pixel intensities in the original image of the surface, which corresponds roughly to the strength of the diffuse component of reflectance (ρ_d). Most illumination maps contain regions of low illumination, where the specular component contributes little to observed radiance. The darkest areas of the observed sphere therefore prove indicative of its diffuse reflectance. The classifier's second statistic, on the vertical axis of Figure 7, is the variance of a particular wavelet subband. This variance provides a measure of the power in an oriented frequency band in the unwrapped image. Surfaces with brighter, sharper specular components (high ρ_s and low α) tend to score higher on this axis.

When classification based on these two statistics is tested using the cross-validation techniques described above, 6 of 54 samples are classified incorrectly. Even the classifier of Figure 7, trained on all 54 samples, misclassifies three of these samples. Two misclassified images are shown in the figure, together with the closest correctly classified images in the space defined by the two statistics. Although such confusions might be expected from a classifier using only two statistics per image, humans could classify these images correctly, largely by noting the distinctness of the sharpest edges in the misclassified examples.

Adding statistics based on the kurtoses of distributions of wavelet coefficients, a rough measure of the prevalence of sharp edges in the image, substantially improves classifier performance. Including ratios of variances of wavelet coefficients at different scales also helps. Through experimentation on several image sets, we settled on a classifier using 6 statistics: the mean and 10th percentile of the original unwrapped image, the variance of coefficients in the finest and second-finest radially (vertically) oriented subbands, the ratio of these two variances, and the kurtosis of the second-finest radially oriented subband. Cross-validation using this choice of statistics gives an accuracy of 53 out of 54 correct (98%).

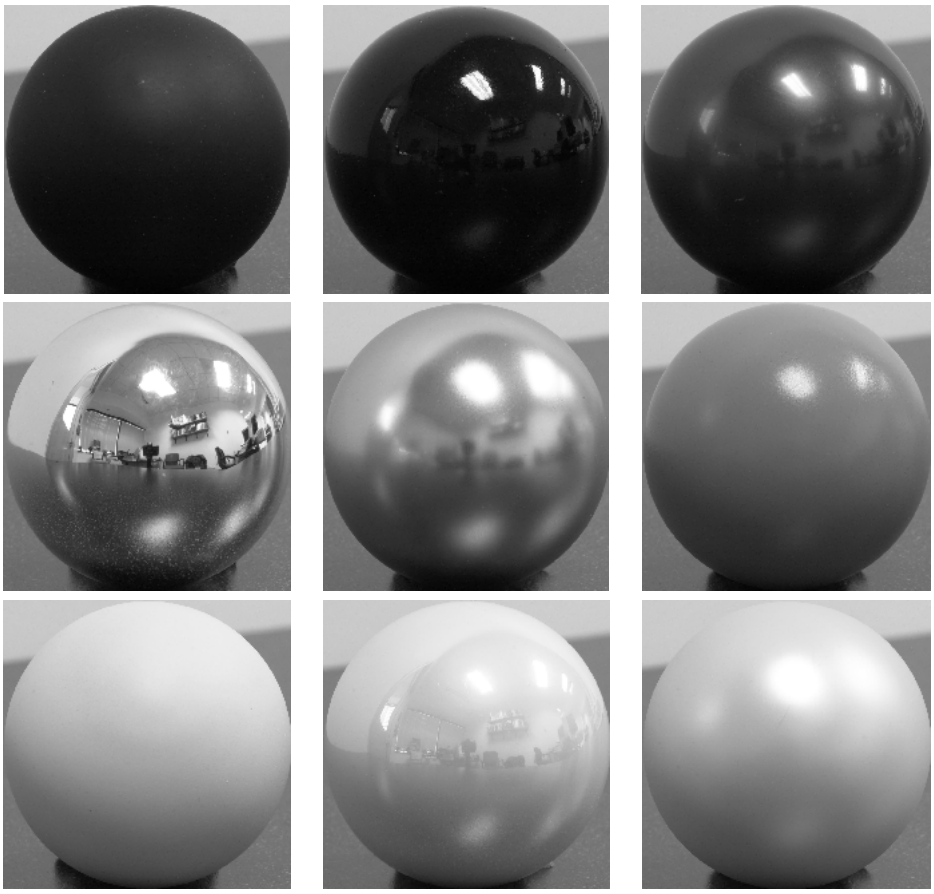

Figure 8. Photographs of the 9 spheres used for the reflectance classification experiment, all under the same illumination. The classification algorithm uses image data only from the surface itself, not from the surrounding background.

5.2. Photographs

We trained and tested a second set of classifiers using images of spheres captured with a Nikon D1 digital camera. Spheres of 9 different materials were photographed under 7 diverse illumination conditions, including both indoor and outdoor settings. Figure 8 shows the 9 spheres under one particular illumination. Images were acquired in 24-bit RGB format and then converted to 8-bit gray scale images for further processing.

Cross-validation of classifiers based on the set of six statistics chosen in Section 5.1 yielded 59 correctly classified samples out of 63 images (94% accuracy). An uninformed random choice in such a 9-way classification task would be 11% accurate. Classification based only on the image mean, normalized for overall strength of illumination, yielded 44% accuracy. Given that each classifier was trained using photographs from only 6 illuminations, 94% accuracy is surprisingly high. In fact, one of the authors misclassified more images than the algorithm.

Figure 9 summarizes the illumination conditions used for each set of photographs by displaying chrome spheres under each illumination. Only one illumination (E) was created by the photographer particularly for the purpose of collecting these photographs. Interestingly, illumination E accounted for 2 of the 4 errors in the cross-validation process. Images under this illumination also proved most difficult for humans to classify, lending credence to the claim that the statistical regularity of "natural" illumination plays an essential role in reflectance estimation.

6. CONCLUSION

This paper formulates the problem of reflectance estimation under unknown illumination, presents a framework for solving this problem using the statistical regularity of real-world illumination, and illustrates preliminary results for

A B C D

E F G

Figure 9. Chrome spheres representing the 7 illumination conditions under which each of the 9 spheres was photographed.

both synthetic images and photographs. Further work will improve the performance of our reflectance estimator and apply it in a more practical setting. First, we plan to apply our estimator to surfaces other than spheres. Our method generalizes naturally to surfaces of other convex geometries, because under the assumption of distant illumination, the observed brightness of a surface patch depends only on its orientation relative to the viewer. Second, we will investigate principled manners to choose the best statistics for classification from the family defined by the distributions of pixel intensities and wavelet coefficients. Third, we hope to utilize additional statistics capturing dependencies of wavelet coefficients across scales, which are known to be significant characterizing of sharp edges and in the modeling of both natural imagery and texture.[19,16]

ACKNOWLEDGMENTS

Marc Talusan and Rebecca Loh provided the photographs of spheres analyzed in Section 5.2. This work was supported by an NDSEG fellowship to R.O.D., by NIH Grant EY11005-04 to E.H.A., by a grant from NTT to E.H.A., and by ONR Grant N00014-00-1-0089 and AFOSR Grant F49620-00-0362 to A.S.W.

REFERENCES

1. R. T. Fleming, R. O. Dror, and E. H. Adelson, "Surface reflectance estimation under unknown natural illumination," in review.
2. S. Tominaga and N. Tanaka, "Estimating reflection parameters from a single color image," *IEEE Computer Graphics and Applications* **20**, pp. 58–66, September/October 2000.
3. Y. Sato, M. D. Wheeler, and K. Ikeuchi, "Object shape and reflectance modeling from observation," in *Proceedings of SIGGRAPH*, 1997.
4. S. R. Marschner, *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, Ithaca, NY, 1998.
5. Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse global illumination: Recovering reflectance models of real scenes from photographs," in *Computer Graphics (SIGGRAPH)*, pp. 215–24, (Los Angeles), 1999.
6. Y. Yu and J. Malik, "Recovering photometric properties of architectural scenes from photographs," in *Computer Graphics (SIGGRAPH)*, pp. 207–17, (Orlando, Florida), 1998.

7. S. Shafer, "Using color to separate reflection components," *Color Research and Application* **10**, pp. 210–218, 1985.
8. A. S. Glassner, *Principles of Digital Image Synthesis*, vol. 2, Morgan Kaufmann, San Francisco, 1995.
9. G. J. Ward, "Measuring and modeling anisotropic reflection," *Computer Graphics (SIGGRAPH)* **26**(2), pp. 265–72, 1992.
10. G. W. Larson and R. Shakespeare, *Rendering with Radiance: The Art and Science of Lighting Visualization*, Morgan Kaugmann, San Francisco, 1998.
11. 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.
12. 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.
13. E. P. Simoncelli and E. H. Adelson, "Subband transforms," in *Subband Image Coding*, J. W. Woods, ed., ch. 4, pp. 143–192, Kluwer Academic Publishers, Norwell, MA, 1990.
14. D. L. Ruderman, "The statistics of natural images," *Network: Comput. Neural Syst.* **5**, pp. 517–48, 1994.
15. J. Huang and D. Mumford, "Statistics of natural images and models," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 541–47, 1999.
16. E. Simoncelli, "Modeling the joint statistics of images in the wavelet domain," in *Proc SPIE, 44th Annual Meeting*, vol. 3813, (Denver), July 1999.
17. W. Freeman, "Exploiting the generic viewpoint assumption," *IJCV* **20**(3), p. 243, 1996.
18. D. J. Heeger and J. R. Bergen, "Pyramid-based texure analysis/synthesis," *Computer Graphics (SIGGRAPH)* , 1995.
19. J. S. De Bonet, "Multiresolution sampling procedure for analysis and synthesis of texture images," *Computer Graphics (SIGGRAPH)* , pp. 361–368, 1997.
20. P. Schröder and W. Sweldens, "Spherical wavelets: Efficiently representing functions on the sphere," *Computer Graphics (SIGGRAPH)* , 1995.
21. T. Evgeniou, M. Pontil, and T. Poggio, "Regularization networks and support vector machines," *Advances in Computational Mathematics* **13**(1), pp. 1–50, 2000.
22. R. Collobert and S. Bengio, "Support vector machines for large-scale regression problems," Tech. Rep. IDIAP-RR 00-17, IDIAP, Martigny, Switzerland, 2000.
23. E. H. Adelson, "Lightness perception and lightness illusions," in *The Cognitive Neurosciences*, M. Gazzaniga, ed., pp. 339–351, MIT Press, Cambridge, MA, 1999.
24. J. R. Boyack and A. K. Juenger, "Brightness adjustment of images using digital scene analysis." U.S. Patent US5724456, March 1998.