



massachusetts institute of technology — artificial intelligence laboratory

Recognition of Surface Reflectance Properties from a Single Image under Unknown Real-World Illumination

Ron O. Dror, Edward H. Adelson and
Alan S. Willsky

AI Memo 2001-033

November 2001

Abstract

This paper describes a machine vision system that classifies reflectance properties of surfaces such as metal, plastic, or paper, under unknown real-world illumination. We demonstrate performance of our algorithm for surfaces of arbitrary geometry. Reflectance estimation under arbitrary omnidirectional illumination proves highly underconstrained. Our reflectance estimation algorithm succeeds by learning relationships between surface reflectance and certain statistics computed from an observed image, which depend on statistical regularities in the spatial structure of real-world illumination. Although the algorithm assumes known geometry, its statistical nature makes it robust to inaccurate geometry estimates.

1. Introduction

Humans recognize objects visually on the basis of material composition as well as shape. A person would rarely confuse a silver knife with a disposable plastic knife. Figure 1 shows pedestal-shaped objects with different reflectance properties, each imaged in two different real-world settings. The two images of each pedestal are completely different at the pixel level because illumination varies from one location to another. A human observer not only recognizes the similarity between the two surfaces of a given reflectance, but can also classify the images in the third column, which have novel geometry and illumination, into one of the four reflectance categories. This paper describes a computer algorithm with a similar ability to recognize surface reflectance. We build on our previous work [7] by extending the algorithm to surfaces of arbitrary known geometry and by dealing with irregular sampling of surface normals.

The reflectance estimation problem proves underconstrained in the absence of restrictions on illumination. The apparently matte sphere in Figure 1 could be a perfect chrome reflector; a chrome sphere simply reflects its environment, so it could in principle take on an arbitrary appearance. Real-world illumination is highly variable, with direct or reflected light incident on a surface from nearly every direction. Yet in a statistical sense, illumination is far from arbitrary. A chrome surface typically reflects sharp edges, bright light sources, and other common features of the visual world. The regular spatial structure of real-world illumination translates to recognizable characteristics in the appearance of different surfaces.

Our primary goal in studying reflectance is to build a machine vision system capable of recognizing surface materials and their properties. We wish to identify a surface as plastic, metal, or paper; to tell whether a surface is wet or




Figure 1. The task addressed by our classifier. Using images of several surface materials under various illuminations as a training set, we wish to classify novel objects under novel illumination according to their surface material.

dry, dirty or clean. Other applications further motivate the reflectance estimation problem. Reconstruction of a scene from photographs for computer graphics requires inference of both the geometry and the reflectance of visible surfaces. An ability to estimate reflectance from image data under unknown lighting conditions may help overcome the limitations of shape-from-shading algorithms that assume reflectance is known in advance, and of classical algorithms for motion or stereo estimation that assume Lambertian surface reflectance.

The importance of reflectance models in computer graphics has motivated several researchers to develop image-based reflectance estimation techniques. Many of these assume point source illumination [25, 21, 16] and therefore do not apply to photographs taken in the natural world under complex, unknown lighting conditions. Yu *et al.* [29] and Boivin and Galalowicz [2] iteratively estimate both the illumination and reflectance of every surface patch in a scene. To ensure that their optimization will converge to a unique and correct solution, they require a complete geometric model of the surrounding scene, a reasonable initial estimate for illumination, and either multiple photographs

or human interaction in the estimation process. 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 chose to base our classifier on only a single monochrome image of the surface of interest, because we wish to determine what information the basic image structure captures about reflectance. Furthermore, we use only the portion of the image corresponding to the surface itself. We have found that humans can estimate certain surface reflectance properties even in the absence of these cues [11], although their estimates may be biased by variation in illumination [11] or surface geometry [17]. In the future, we hope to improve our results by drawing on additional sources of information, including color spectral decompositions [22, 25], motion cues, and visual context.

We assume in this paper that the surface under observation has homogeneous reflectance properties. While we allow arbitrary surface geometry, we assume that geometry is known in advance. Preliminary results indicate that our algorithm performs robustly even when assumed geometry differs from actual geometry (Section 4.3).

2. Problem formulation

2.1. BRDF estimation

To predict the appearance of an opaque surface patch under any pattern of illumination, one needs the bidirectional reflectance distribution function (BRDF) of the surface patch. The BRDF, a mathematical definition of reflectance, specifies what proportion of the light incident from each possible illumination direction is reflected in each possible view direction. It is a function of two directions in the three-dimensional world, and therefore of four angular variables. One can compute the brightness of a surface patch viewed from a particular direction by performing a weighted integral over illumination from all directions, with the BRDF specifying the weighting [12].

Rendering an image given complete knowledge of surface reflectance and illumination is computationally expensive but straightforward. Recovering a surface BRDF from an observed image under unspecified illumination, on the other hand, is a highly underconstrained problem. The image data is a function of two variables, while the BRDF is a function of four. Moreover, illumination from every direction is unknown and can vary across the surface. The BRDF must conserve energy and satisfy a symmetry property known as reciprocity [12], but the space of possible BRDFs remains huge. Ramamoorthi and Hanrahan [20] have shown that even when one is given images of a homogeneous surface from all possible view directions, dif-

ferent combinations of illumination and reflectance can explain the observations.

We address a problem that is more tractable than general BRDF estimation, but that remains underconstrained. We classify surface reflectance, attempting to select one of a finite set of reflectances that most closely represents the reflectance of an observed surface. The candidate reflectances may be specified by different parameter settings of one of the reflectance models common in computer graphics, such as the Phong [18] or Ward [26] models. Alternatively, the reflectances can be arbitrary BRDFs, specified as lookup tables or by a set of photographs of a real surface. Figure 1 illustrates one instantiation of this problem.

2.2. Prior information on illumination

One could completely describe illumination at a particular point in the world as a panoramic image specifying the light a camera would see as it looked from that point in every direction. By combining multiple photographs taken from a single point, one can compose such a complete spherical image. These images, known as “environment maps” or “illumination maps,” are used in graphics to render an object as it would appear at a particular point in space [5]. The underlying assumption is that all sources of direct and indirect illumination are sufficiently distant that the illumination map will change little as one moves across the object surface.

One might think of reflectance estimation as a system identification problem, where the BRDF represents unknown system parameters of interest, and the observed image of a surface represents the system’s output signal. The problem is difficult because the illumination, the input signal, is unknown. The problem becomes tractable if the input signal has predictable statistics. Although illumination maps differs in field of view and dynamic range from the photographs studied in the natural image statistics literature [10, 14], they share many of the statistical regularities of typical photographs [8]. For example, marginal and joint distributions of wavelet coefficients at various scales and orientations exhibit similar heavy-tailed distributions from image to image. Figure 2 shows an example.

Figure 3 shows synthetic images of two identical spheres under different illuminations. Humans identify surface reflectance more easily in image B, rendered under a photographically-acquired illumination map, than in image A, rendered under point source illumination. We found a similar effect when comparing photographs of a sphere in a normally illuminated room and in a black room with a point light source. Point source illumination does not share the statistics of typical natural illumination. Previous work on reflectance estimation has often considered the case of point source illumination as a convenient starting point. We




Figure 2. Distributions of wavelet coefficients at three successive scales for an indoor illumination map (left) and an outdoor illumination map (right). The distributions have high kurtosis, with variance increasing at coarser scales. Wavelet coefficients were computed with a nine-tap quadrature mirror filter pyramid [23] from log-luminance equal-area cylindrical projections of the spherical maps.




Figure 3. (A) A sphere rendered under point source illumination. (B) The same sphere rendered under photographically-acquired illumination.

wish instead to take advantage of the statistical complexity of natural illumination in estimating reflectance.

3. A method for reflectance classification

3.1. 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:

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

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

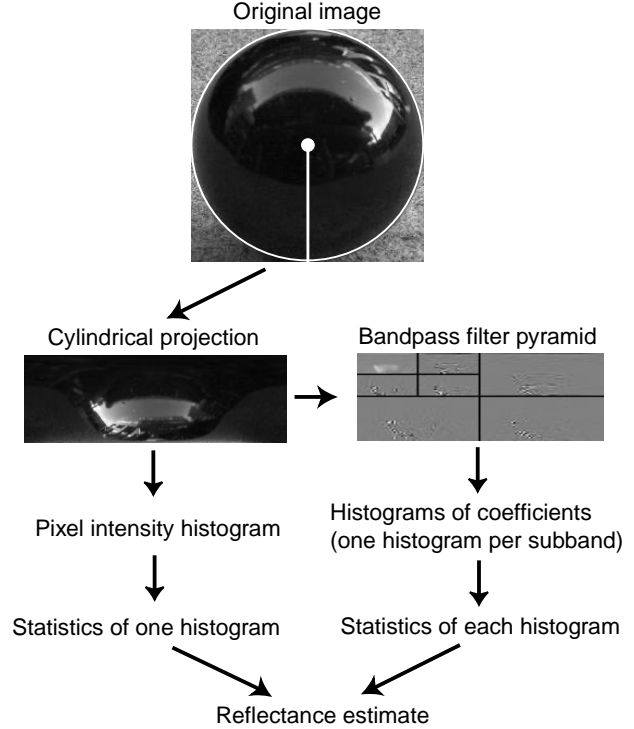


Figure 4. Flowchart for computation of image features for classification, illustrated for spherical surface geometry. Sections 3.3 and 3.5 discuss modifications for arbitrary surface geometry.

prior probability of any given illumination map $P(I)$ explicitly, integration over all possible illuminations is computationally daunting. We developed an alternative technique for practical reflectance estimation.

3.2. A machine learning approach

We apply machine learning techniques to determine relationships between surface reflectance and statistics of the observed image. Our choice of statistics is inspired by the natural image statistics and texture analysis literatures [8, 10, 14, 13, 19]. Figure 4 illustrates our approach in the case where the observed surface is spherical. We first project the observed data, defined as a function of surface orientation, onto a plane. Next, we compute a set of statistics on the pixel intensity distribution of the image itself and on distributions of outputs of a set of oriented band-pass filters applied to the image. We use either photographs or synthetic images rendered under photographically-acquired illumination to train a reflectance classifier based on these statistical features.




Figure 5. Images of surfaces of three geometries, each mapped to an equivalent sphere of the same reflectance and illumination, then “unwrapped” to a cylindrical projection of the sphere. All three original surfaces have the same reflectance and were rendered under the same illumination. Black regions within the equivalent spheres and cylindrical projections indicate missing data. The visible portion of the pedestal-shaped object lacks surface normals in many directions. The worm-shaped surface has all the normal directions of a sphere, but they are sparsely sampled due to high surface curvature.

3.3. Accounting for surface geometry

When a distant viewer observes a convex surface under distant illumination, the brightness of a surface point depends only on local surface orientation. We wish to analyze the image data as a function of surface orientation rather than spatial position. We therefore map the brightness of each observed surface point to the point with the same normal on a hypothetical sphere with the same reflectance and illumination. Figure 5 shows several examples in which we partially reconstruct the appearance of such a sphere from an observed image. Because some surface normals are sparsely sampled or absent in the observed image, some points on the surface of the equivalent sphere are not recovered. Section 3.5 describes our method for handling this missing data.

One would ideally compute all statistics directly in the spherical domain. However, methods for natural image and texture analysis have been developed in a planar domain. We therefore chose to project the hemispherical data onto a plane before performing further analysis. In order to

preserve stationarity of local image statistics to the extent possible, we perform a cylindrical projection of the visible hemisphere. Intuitively, this amounts to “unwrapping” an orthographic projection of the sphere about its center, as shown in Figure 4. We experimented with a number of cartographic projections for this purpose and found that the choice of projection has only a minor impact on classifier performance. Figure 5 shows projections for several surfaces of different geometries but identical reflectance and illumination. The region near the center of the sphere is expanded significantly in the projection process. Because we map pixels from the original image to pixels in the cylindrical projection, the corresponding region of the projection lacks information for many surface normals and therefore contributes less to the computation of image statistics (Section 3.5).

3.4. Choice of features

Our choice of statistical features for classification was influenced by the texture analysis work of Heeger and Bergen [13], and by Nishida and Shinya’s finding that luminance histograms of observed images influence human perception of surface reflectance [17]. First, we construct a histogram of pixel intensities in the original observed image. Next, we filter the image using vertical and horizontal derivatives of two-dimensional Gaussians. We construct a histogram to approximate the distribution of filter outputs at each scale and orientation. Our experience is that the precise choice of a multiscale decomposition is not critical to classifier performance.¹

We compute a set of numerical statistics to characterize each distribution. These include the mean, variance, skew, kurtosis, and the 10th, 50th, and 90th percentiles. Variances of the filter outputs at different scales provide an approximate measure of the spectral power in different frequency bands. Kurtoses of these distributions provide a rough measure of image edginess, because edges tend to produce extreme bandpass filter outputs, leading to heavy-tailed (highly kurtotic) distributions.

Figure 6 shows how images of surfaces of the same reflectance cluster in feature space when observed under different natural illuminations. For purposes of visualization, we have chosen only two features, both of which have intuitive significance. The horizontal axis denotes the 10th percentile of the distribution of pixel intensities in the original image. This corresponds roughly to the strength of the diffuse component of reflectance. Most illumination maps contain regions of low illumination, where the specular component contributes little to observed radiance. The

¹We originally worked with a pyramidal decomposition based on nine-tap symmetric quadrature mirror filters [23]. We currently use derivatives of Gaussians because they lend themselves to the normalized differential convolution approach described in Section 3.5.




Figure 6. Solid symbols indicate locations in a two-dimensional feature space of images of spheres of six reflectances, each rendered under nine different real-world illuminations. Lines separate the regions that the SVM classifier (Section 3.6) assigns to different reflectances.

darkest areas of an observed surface therefore prove indicative of its diffuse reflectance. The classifier’s second statistic, on the vertical axis of Figure 6, is the variance of a particular bandpass filter output. Surfaces with brighter, sharper specular components tend to score higher on this axis. Classifiers based on more statistics prove more accurate.

3.5. Feature computation with sparse data

The computation of features based on filter outputs is complicated by the missing samples in images to be analyzed. The reconstructed equivalent spheres (Figure 5) may lack entire regions of data because the corresponding surface normals are not present on the observed surface. Other areas of the reconstructed spheres contain small gaps due to sparse sampling of surface normals in the original image. We estimate bandpass filter outputs in these regions using an instantiation of normalized differential convolution [28]. At each point in the projected image, we find the linear combination of a constant function, a horizontally-oriented ramp, and a vertically-oriented ramp that best approximate, in a weighted mean-squared sense, the observed samples over a Gaussian window surrounding that point. The coefficients of the two ramp functions represent our estimate of the derivative-of-Gaussian filter outputs at that point. The size of the Gaussian window varies by a factor of two from scale to scale. Following Westelius [27], we compute a confidence measure for the filter output estimates based on the

amount of data present within a particular Gaussian window. When computing a histogram for filter outputs at each scale, we use only those outputs whose associated confidence exceeds a threshold value.

3.6. Classification techniques

In principle, adding more features to the classifier inputs can only provide additional information about reflectance. In practice, we have only a limited amount of training data, because either the surface images themselves or the illumination maps used to render them must be acquired photographically in real-world environments. An excessive number of features will therefore lead to degradation in classifier performance. The results illustrated in this paper are based on a set of hand-selected features: the mean and 10th percentile of the pixel intensity distribution, the variance of the horizontally and vertically oriented filter outputs at the three finest scales, and the kurtosis of vertically oriented filter outputs at the finest scale. A separate paper [7] compares performance with various feature sets and describes an automatic method for feature selection.

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 [9]. Our implementation utilizes the SVMTool software [4] with Gaussian kernels to train and apply SVM classifiers. SVMTool uses a one-against-all voting scheme to perform multiclass classification. In [7], we compared the performance of various classifier and found that SVMs substantially outperform nearest neighbor and k -nearest neighbor classifiers when the number of features is large compared to the number of training images. Figure 6 shows the class boundaries determined by an SVM classifier based on two statistics.

3.7. Lightness ambiguity

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 [1, 3], 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 near the surface under observation.

4. Results

4.1. Image sets

We trained and tested classifiers on photographs as well as on synthetic images rendered under photographically-acquired real-world illumination maps. We photographed spheres of nine different materials under seven diverse illumination conditions indoors and outdoors using a Nikon D1 digital camera. Figure 7 shows examples of these 8-bit gray-scale images. The full image set is available at http://www.ai.mit.edu/people/rondror/sphere_photos/.

To create synthetic images, we used Ward’s *Radiance* package [15], which efficiently implements the Ward reflectance model [26], a physically realizable variant of the Phong shading model.² Our rendering methodology is similar to that of Debevec [5]. For nonconvex objects such as those in Figures 1 and 8, our rendering includes the effects of interreflection. Our classifier uses only the final rendered images for training and testing.

We chose a set of six reflectances specified by Ward model parameters corresponding to common materials of distinctly different appearances (Figure 8). We rendered objects with each of these reflectances and with several geometries under two sets of photographically-acquired illumination conditions. The first set consisted of nine high dynamic range spherical illumination maps from Debevec’s Light Probe Image Gallery [6] (<http://www.debevec.org/Probes/>), which represent diverse lighting conditions from four indoor settings and five outdoor settings. The second set included 96 illumination maps based on high dynamic range imagery acquired by Teller *et al.* [24] in outdoor urban environments (<http://city.lcs.mit.edu/data>). The images in Figure 1 were rendered under one of Debevec’s illumination maps, while those in Figure 8 were rendered under one of Teller’s.

4.2. Performance

We trained reflectance classifiers on images of spheres, and then tested them on images of surfaces with various geometries. When testing classification accuracy, we avoided using any one illumination for both training and testing. In a variant of leave-one-out cross-validation, we tested a

²The isotropic Ward reflectance model takes the form

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

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.




Figure 7. Photographs of nine spheres in a single location. Each sphere was photographed in seven locations.

classifier on images corresponding to one illumination after training it on images corresponding to the remaining illuminations. By leaving out each illumination and repeating this process, we were able to use images rendered under all illuminations as test images.

Our classifier achieved 94% accuracy for the photographic data set illustrated in Figure 7, misclassifying 4 of 63 images. A naive random choice in such a 9-way classification task yields 11% accuracy, while a classifier based only on mean surface brightness, normalized for overall strength of illumination, yields 44% accuracy. An accuracy of 94% rivals human performance; one of the authors misclassified more images in an informal test. The classifier achieved these results even though it was trained on images under only 6 illuminations, and even though the compressive nonlinearity applied by the digital camera produces images that are not linear in luminance. Two of the four misclassifications corresponded to the one illumination created by the photographer particularly for the purpose of collecting these photographs. These images, photographed directly under a desk lamp in an otherwise dark room, also proved most difficult for humans to classify.

We trained and tested classifiers separately for the 6 Debevec illuminations and the 96 Teller illuminations, in each case using the previously described variant of leave-one-out cross-validation. When trained and tested on spheres, our classifiers achieved 98% accuracy for the Debevec illuminations, and 99.5% accuracy for the larger set of Teller illuminations. For comparison, classifiers trained on mean




Figure 8. Synthetic images of worm-shaped surfaces of 6 different reflectances, each rendered under one of Teller’s illumination maps. Ward model parameters are as follows: black matte, $\rho_d = .1$, $\rho_s = 0$; black shiny, $\rho_d = .1$, $\rho_s = .1$, $\alpha = .01$; chrome, $\rho_d = 0$, $\rho_s = .75$, $\alpha = 0$; gray shiny, $\rho_d = .25$, $\rho_s = .05$, $\alpha = .01$; white matte, $\rho_d = .9$, $\rho_s = 0$; white shiny, $\rho_d = .7$, $\rho_s = .25$, $\alpha = .01$.

brightness alone were respectively 52% and 30% accurate.

Next, we trained the classifiers on images of spheres and tested them on images of the “pedestal” (Figure 1) and the “worm” (Figure 8). The classifiers maintained high performance. For the Debevec illuminations, we obtained 98% accuracy when testing on the pedestal shape, and 94% when testing on the worm. For the Teller illuminations, we obtained 94% accuracy for the pedestals and 97% accuracy for the worms.

Performance tends to be slightly poorer for geometries such as the pedestal and worm for several reasons. First, these shapes incompletely sample the space of surface normals (Figure 5). In the case of the pedestal, certain regions of the equivalent sphere are missing entirely, while in the case of the worm, large regions are sampled sparsely due to the high curvature of the surface. Second, these shapes are not convex. Illumination at two surface points with identical normals may differ because of interreflection or light source occlusion. These differences may introduce artifacts in the reconstructed image of the equivalent sphere. Performance for non-convex geometries may improve if the surface is first split into regions such that few points in the same region share the same surface normal. One could compute bandpass filter outputs for projections of each of these regions separately, then pool filter outputs from different regions before computing their distribution statistics.




Figure 9. Image of a sphere (left) and an ellipsoid (right) with identical reflectance rendered under identical illumination. The major axis of the ellipsoid, directed toward the viewer, is twice the length of its other two axes.

4.3. Robustness to inaccurate geometry

In contrast to methods that rely on explicit inverse rendering of observed images, our reflectance estimation method is robust to some errors in the assumed geometry. Figure 9 shows a sphere and an elongated ellipsoid of identical reflectance under identical illumination. Although both surfaces have circular contours, the images differ significantly at the pixel level. We applied the six-reflectance classifier to images of the ellipsoid rendered under the 9 Debevec illuminations, assuming incorrectly that its geometry was spherical. The classifier correctly labeled 51 of 54 images (94% accuracy), compared to 53 of 54 when the assumed geometry was accurate. A human observer, likewise, easily recognizes the similarity in reflectance of the two objects, while hardly noticing the difference in geometry.

5. Conclusions

This paper demonstrates a method for reflectance classification from single monochrome images that applies to surfaces of arbitrary geometry in unknown real-world scenes. Our solution to this otherwise ill-posed problem relies on the predictable statistical structure of real-world illumination. The statistical nature of our approach confers a degree of robustness to incorrect geometry estimates. We hope to use these techniques as part of a more general system for visual recognition of materials.

Acknowledgments

Marc Talusan and Rebecca Loh carried out the photography of the spheres analyzed in Section 4. 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. Lightness perception and lightness illusions. In M. Gazzaniga, editor, *The Cognitive Neurosciences*, pages 339–351. MIT Press, Cambridge, MA, 1999.
- [2] S. Boivin and A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. *Computer Graphics (SIGGRAPH)*, 2001.
- [3] J. R. Boyack and A. K. Juenger. Brightness adjustment of images using digital scene analysis. U.S. Patent US5724456, March 1998.
- [4] R. Collobert and S. Bengio. Support vector machines for large-scale regression problems. Technical Report IDIAP-RR 00-17, IDIAP, Martigny, Switzerland, 2000.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] R. O. Dror, T. Leung, A. S. Willsky, and E. H. Adelson. Statistics of real-world illumination. In *Proceedings of CVPR*, 2001.
- [9] T. Evgeniou, M. Pontil, and T. Poggio. Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13(1):1–50, 2000.
- [10] 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.
- [11] 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.
- [12] A. S. Glassner. *Principles of Digital Image Synthesis*, volume 2. Morgan Kaufmann, San Francisco, 1995.
- [13] D. J. Heeger and J. R. Bergen. Pyramid-based texture analysis/synthesis. *Computer Graphics (SIGGRAPH)*, 1995.
- [14] J. Huang and D. Mumford. Statistics of natural images and models. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 541–47, 1999.
- [15] G. W. Larson and R. Shakespeare. *Rendering with Radiance: The Art and Science of Lighting Visualization*. Morgan Kaufmann, San Francisco, 1998.
- [16] S. R. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, Ithaca, NY, 1998.
- [17] S. Nishida and M. Shinya. Use of image-based information in judgments of surface-reflectance properties. *J. Opt. Soc. Am. A*, 15:2951–2965, 1998.
- [18] B.-T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–7, 1975.
- [19] 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.
- [20] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. *Computer Graphics (SIGGRAPH)*, 2001.
- [21] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *Computer Graphics (SIGGRAPH)*, 1997.
- [22] S. Shafer. Using color to separate reflection components. *Color Research and Application*, 10:210–218, 1985.
- [23] E. P. Simoncelli and E. H. Adelson. Subband transforms. In J. W. Woods, editor, *Subband Image Coding*, chapter 4, pages 143–192. Kluwer Academic Publishers, Norwell, MA, 1990.
- [24] 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.
- [25] S. Tominaga and N. Tanaka. Estimating reflection parameters from a single color image. *IEEE Computer Graphics and Applications*, 20:58–66, September/October 2000.
- [26] G. J. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics (SIGGRAPH)*, 26(2):265–72, 1992.
- [27] C.-J. Westelius. *Focus of Attention and Gaze Control for Robot Vision*. PhD thesis, Linköping University, Linköping, Sweden, 1995.
- [28] C.-F. Westin, K. Nordberg, and H. Knutsson. On the equivalence of normalized convolution and normalized differential convolution. In *ICASSP*, volume V, pages 457–460, 1994.
- [29] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Computer Graphics (SIGGRAPH)*, pages 215–24, Los Angeles, 1999.