Reflectance and Natural Illumination from a Single Image

Stephen Lombardi and Ko Nishino

Department of Computer Science
Drexel University, Philadelphia, PA 19104, USA
{sal64,kon}@drexel.edu

Abstract. Estimating reflectance and natural illumination from a single image of an object of known shape is a challenging task due to the ambiguities between reflectance and illumination. Although there is an inherent limitation in what can be recovered as the reflectance band-limits the illumination, explicitly estimating both is desirable for many computer vision applications. Achieving this estimation requires that we derive and impose strong constraints on both variables. We introduce a probabilistic formulation that seamlessly incorporates such constraints as priors to arrive at the maximum a posteriori estimates of reflectance and natural illumination. We begin by showing that reflectance modulates the natural illumination in a way that increases its entropy. Based on this observation, we impose a prior on the illumination that favors lower entropy while conforming to natural image statistics. We also impose a prior on the reflectance based on the directional statistics BRDF model that constrains the estimate to lie within the bounds and variability of real-world materials. Experimental results on a number of synthetic and real images show that the method is able to achieve accurate joint estimation for different combinations of materials and lighting.

1 Introduction

The appearance of an object is determined by surface geometry, reflectance, and illumination. Decomposing an image into these physical entities is a challenging task that underlies many of the longstanding computer vision problems. Among the many instances of this decomposition problem, estimating the reflectance of an object of known shape but taken under unknown illumination is particularly important. Solving this key problem can provide significant information that can aid in recognizing materials and objects in a scene. In particular, achieving this dual estimation from images taken under natural illumination is of paramount importance for scene understanding in the real world. At first glance, stepping out from the conventional dark-room setup of a single or multiple point light sources, however, seems to add overwhelming complexity to the problem. Natural illumination, just like natural images, embodies structures in its spatial layout and color distribution. In this paper, we are interested in whether such inherent structures can play to our side. We show that when exploited together with the intrinsic structures found in the reflectance of real-world materials, it enables robust simultaneous estimation of reflectance and natural illumination from a single image.

The difficulty of estimating both the reflectance and illumination is exacerbated by the inherent ambiguities between the two variables. An obvious ambiguity arises from

A. Fitzgibbon et al. (Eds.): ECCV 2012, Part VI, LNCS 7577, pp. 582–595, 2012.

the color constancy problem—an image of a red object can be explained with either a white reflectance and red illumination or red reflectance and white illumination. A more complex ambiguity lies in the frequency composition of reflectance and illumination. The reflectance acts as a bandpass filter by attenuating arbitrary frequency components of the illumination [1]. Real-world reflectance may take on various frequency spectra, but in general they act as a lowpass filter. Because of this, one can only hope to recover illumination up to a certain frequency. This problem applies in reverse as well—low frequency illumination conceals higher frequency features of the reflectance.

Previous methods have typically made strong, limiting assumptions about either the illumination (e.g., point light source(s)) [2–5] or the reflectance (e.g., Lambertian) [6–9] that undermines their practicality. The recent work by Romeiro and Zickler [10], named "blind reflectometry," achieves reflectance estimation under natural illumination by marginalizing over a distribution of possible lighting environments to overcome the inherent ambiguity between reflectance and illumination. Although this results in a reasonably accurate estimation of reflectance, the method cannot explicitly recover the illumination. Furthermore, to circumvent the color constancy problem, the method only estimates a monochrome reflectance. This leads to serious limitations in predicting the appearance of objects, such as incorrect colors in highlights.

Our goal is to jointly estimate *both* the reflectance and natural illumination in full color. We believe it is essential to explicitly estimate both to enable a host of applications that would benefit from properly predicting object appearance. For instance, in object recognition and tracking, the task would become easier if we were able to accurately predict how the target object would look from a different viewpoint given a single image of it. This can be achieved only when both the reflectance and the illumination are explicitly recovered. This is true even though one of them, usually the illumination estimate, is band-limited in frequency as those explicit estimates are sufficient to compute the object appearance in any other pose under that illumination. Explicit estimates would also enable accurate appearance prediction of objects with reflectance that have lower frequencies or under lower frequency illumination, too. To this end, our goal is not only to perform blind reflectometry but also *blind light probe*.

In this paper, we introduce a novel method for jointly estimating the illumination of a natural scene and the reflectance of an object with homogeneous material from a single image. Resolving the ambiguities and overcoming the loss of information requires that we make strong and accurate assumptions about typical illumination environments and reflectance properties. For this, we derive data-driven and information theoretic constraints that faithfully encode the variability and interaction of real-world reflectance and natural illumination. We then introduce a probabilistic formulation that enables the seamless incorporation of these assumptions as prior distributions on the latent variables to jointly estimate the most "realistic" reflectance and illumination.

To constrain the reflectance estimate to lie within the space spanned by those of real-world materials, we build a statistical prior using the directional statistics bidirectional reflectance distribution function (DSBRDF) model [11]. Based on the functional principal component analysis that extracts a linear subspace of real-world reflectance [12], we impose a tight yet analytically simple statistical prior that faithfully encodes the variability of real-world reflectance. To tame the natural illumination estimate, we

consider two constraints. We model the illumination as a discrete image whose axes are inclination and azimuth angles in the spherical coordinate system (i.e., latitude and longitude). This allows us to view natural illumination as a wide-angle view natural image and directly leverage natural image statistics—a heavy-tail gradient distribution—to constrain the estimate. We also observe and empirically confirm that the reflectance modulates the natural illumination in a way that increases its entropy. The intuition is that the reflectance acts as a blur kernel on the natural illumination and thus smooths out the histogram of its color values leading to increased entropy. To this end, we constrain the estimation to favor lower entropy illumination estimates through a prior on the illumination. These priors are incorporated into our Bayesian formulation that we solve for the maximum a posteriori (MAP) estimate. Joint estimation of reflectance and illumination is achieved through expectation maximization (EM) iteration whose maximization step involves alternating between the estimations of reflectance and illumination separately.

We demonstrate the effectiveness of our method on a number of synthetic and real images. The experimental evaluation on synthetic images include a large number of combinations of different reflectance and natural illuminations to empirically study the stability of the estimation on different types (frequency characteristics) of those two variables. The results clearly demonstrate that our method can explicitly recover both the reflectance and illumination, up to the inherent limitation, in full color. As far as we know, our work is the first to demonstrate such capability.

2 Representing and Constraining Natural Illumination

To achieve accurate joint estimation of reflectance and natural illumination, we should fully leverage the rich information imparted by natural illumination onto the single image at hand. We can think of arbitrary natural illumination as an infinite set of point lights. This interpretation suggests that we are getting in a single image the information typically only available in multiple images. We are capturing richer information regarding the reflectance (i.e., more angular samples), but it is multiplexed into a single color vector at each surface point. The challenge then lies in demultiplexing these observations into the true natural illumination and reflectance. This can be done robustly, if we could guide the estimation to identify the most likely illumination out of many potential combinations of illumination and reflectance.

We first need to devise an appropriate representation of illumination that allows us to recover as much detail as is theoretically possible. To facilitate this, we use a spherical panorama by discretizing the incident illumination sphere. This representation, which we refer to as the illumination map $\mathbf{L}$, is a 2D image whose vertical and horizontal axes are respectively the inclination θ and azimuth angles ϕ in the spherical coordinate system. This representation allows us to control the level of detail we wish to model by modifying the granularity of the discretization. More important, we may treat the illumination as a wide-angle view image on which we may place image-based priors. We describe two such priors: a natural image statistics prior and our novel entropy-based prior.

2.1 Illumination as a Natural Image

As the natural illumination that we want to recover is essentially an image of the world surrounding an object captured as a directional environment map at where the object sits and represented with a spherical panorama, we can interpret it as a wide-angle view natural image. This also means that we can leverage well-studied statistics of natural images to properly constrain our natural illumination estimate. Previous works [13] on natural image statistics has found that the gradients of natural images form a distribution with a heavy tail. In particular, we model the heavy-tailed distribution of the gradients in the illumination map with a hyper-Laplacian,

$$p_s(\mathbf{L}) = \frac{1}{Z} \exp\left[-\sum_{\theta,\phi} \sum_{(\theta',\phi') \in N(\theta,\phi)} |\mathbf{L}_{\theta,\phi} - \mathbf{L}_{\theta',\phi'}|^{\alpha} \right], \tag{1}$$

where Z is the partition function that normalizes the distribution, $N(\theta, \phi)$ is the set of neighboring pixels of θ and ϕ in the illumination map, and α is the exponent of the hyper-Laplacian. This distribution encourages the illumination estimate to lie in the space of real-world natural illumination.

2.2 The Entropy Increase by Reflection

Our goal is to properly constrain the space of possible illumination estimates that is made ambiguous during image formation. Recall that part of the reason why this ambiguity exists is that the material reflectance acts on the illumination as a bandpass filter. We empirically show that, due to the band-limited transmittance of incident irradiance, the entropy of the distribution of reflected radiance becomes higher than when there was no bandpass filtering (i.e., the reflectance has all frequencies–perfect mirror reflection).

Entropy has been studied in many areas of computer vision as a useful information theoretic metric to constrain estimation. Alldrin et al. [14] use the entropy of the distribution of albedo values of a textured surface to resolve the generalized bas-relief ambiguity in photometric stereo. They assume that the true distribution of albedo values is sparsely peaked and use entropy to guide the estimation to find the albedo values that conform to this property. We make similar assumptions but on the natural illumination we aim to recover. Finlayson et al. [15] propose an approach to remove shadows by finding a direction in a 2D chromaticity feature space under which illumination is invariant. The correct invariant direction minimizes the entropy of the image formed by projecting the 2D feature points along it. As it so happens, entropy can also be used to order the set of potential illumination maps that we face during joint estimation.

Figure 1 demonstrates the effect of the reflectance on the entropy of the reflected radiance for a variety of materials. As our intuition suggests, the action of the BRDF as a bandpass filter causes a blurring of the illumination and thus a spreading of the histogram. This, in turn, increases the entropy of the reflected radiance. We'd like to recover the true illumination environment and to do this we assume that the entropy increase in the observed image is due entirely to the BRDF. To this end, we constrain the illumination to have minimum entropy, so that the BRDF will be responsible for causing the increase in entropy of the outgoing radiance.

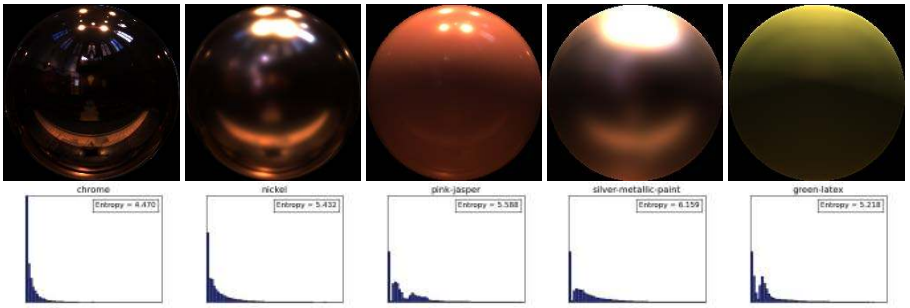

Fig. 1. The BRDF always causes an increase in entropy of the reflected radiance. The most specular materials (e.g., nickel) cause the least increase while the most diffuse (e.g., green-latex) cause the greatest increase. Only a perfect mirror (e.g., chrome), will not increase entropy.

It is with this intuition that we formally derive our entropy prior. Entropy is defined as the expected value of the information context of a random variable. While typically defined for discrete random variables, we use a continuous definition which will allow us to take the derivative with respect to the illumination map later. The entropy of an image is an integral over the distribution of intensity values,

$$H = -\int p(x) \log p(x) dx \,, \tag{2}$$

where in our case $p(x)$ is the histogram of $\mathbf{L}$. To ensure that entropy is differentiable with respect to the illumination map, we write the histogram using kernel density estimation with a Gaussian kernel

$$p(x) = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(x - \mathbf{L}_i)^2}{2\sigma^2} \right] , \tag{3}$$

where N is the number of pixels in the illumination map, $\mathbf{L}_i$ chooses the i^{th} point, and σ^2 is the variance of the Gaussian kernel. Using this expression of entropy, we place a exponential prior on the illumination map with an exponent proportional to the entropy

$$p_e(\mathbf{L}) \propto \exp\left[-H(\mathbf{L}) \right] . \tag{4}$$

3 Modeling the Space of Reflectance

The next problem we face is to properly constrain reflectance. We must take care to choose a model that is flexible enough to represent a large class of reflectance functions but is also amenable to a strong prior. In this paper, we assume that reflectance is well represented by an isotropic Bidirectional Reflectance Distribution Function (BRDF) and so we use the Directional Statistics BRDF (DSBRDF) [11] for the flexibility it provides. Using the DSBRDF model we can write the BRDF as a sum of reflectance lobes

$$\varrho^{(\lambda)}(\theta_d, \theta_h; \kappa^{(\lambda)}, \gamma^{(\lambda)}) = \sum_r \exp\left[\kappa^{(r,\lambda)}(\theta_d)\cos^{\gamma^{(r,\lambda)}(\theta_d)}(\theta_h)\right] - 1\,, \tag{5}$$

where the halfway vector parameterization (i.e., (θ_h, ϕ_h) for the halfway vector and (θ_d, ϕ_d) for the difference vector) [16] is used, $\kappa^{(\lambda)}$ and $\gamma^{(\lambda)}$ are functions that encode the magnitude and acuteness of the reflectance, respectively, of lobe r along the span of θ_d for a particular color channel λ. This allows us to express reflectance functions as a collection of κ and γ curves.

As shown by Nishino and Lombardi [12], these curves can be modeled as a log-linear combination of data-driven basis functions,

$$\kappa^{(r,\lambda)}(\theta_d) = \exp\left[b_\mu(\theta_d; \kappa, r, \lambda) + \sum_i \psi_i b_i(\theta_d; \kappa, r, \lambda)\right], \tag{6}$$

$$\gamma^{(r,\lambda)}(\theta_d) = \exp\left[b_\mu(\theta_d; \gamma, r, \lambda) + \sum_i \psi_i b_i(\theta_d; \gamma, r, \lambda)\right], \tag{7}$$

where b_μ is the mean basis function, b_i is the i^{th} basis function, and ψ_i are the DSBRDF coefficients. We may compute these basis functions from a set of measured reflectance functions using functional principal component analysis (FPCA). This provides the additional benefit of ordering the basis functions by importance so that we can truncate the number of parameters used to achieve a compact representation. Our next step, then, is to model the distribution of the coefficient vector $\mathbf{\Psi}$.

Based on the recent work by Lombardi and Nishino [17], we impose a prior that models the space of real-world isotropic BRDFs using an analytical distribution on $\mathbf{\Psi}$, the projection of measured BRDFs on the DSBRDF basis functions. We adopt a zero-mean multivariate Gaussian prior to model the space of reflectance

$$p(\mathbf{\Psi}) \sim \mathcal{N}(0, \Sigma_\Psi), \tag{8}$$

where Σ_Ψ is computed from the MERL/MIT BRDF database [18]. Note that we exclude the BRDF we try to estimate when conducting the experimental validation using synthetic images.

4 Bayesian Joint Estimation

Our main contribution is the Bayesian formulation that enables the integration of the constraints on the reflectance and illumination as statistical priors based on which we can perform joint estimation. Given an image $\mathbf{I}$, our goal is to find the maximum a posteriori (MAP) estimate of the posterior distribution

$$p(\mathbf{L}, \mathbf{\Psi} \mid \mathbf{I}) \propto p(\mathbf{I} \mid \mathbf{\Psi}, \mathbf{L}) p(\mathbf{\Psi}) p(\mathbf{L})\,. \tag{9}$$

We model image formation as a stochastic process that incurs Gaussian noise. This gives us a likelihood function that is centered on the predicted irradiance $\mathbf{E}$ at pixel $\mathbf{x}$

$$p(\mathbf{I} \mid \mathbf{\Psi}, \mathbf{L}) = \prod_{\mathbf{x}} \mathcal{N}(\mathbf{I}_{\mathbf{x}} \mid \mathbf{E}(\mathbf{\Psi}, \mathbf{L}, \mathbf{N}_{\mathbf{x}}), \sigma_i^2)\,, \tag{10}$$

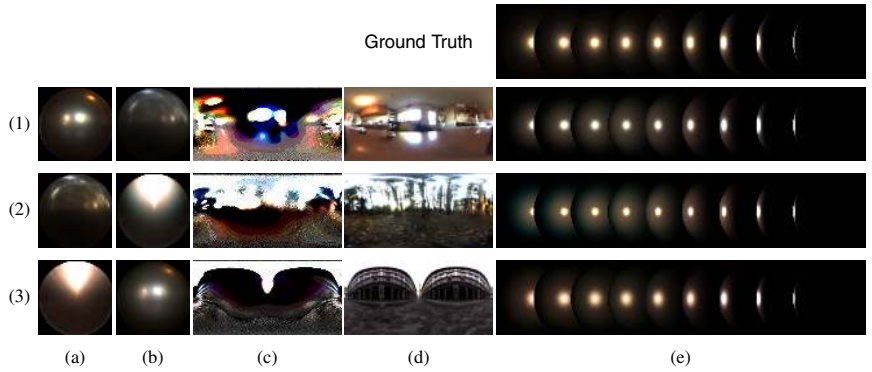

Fig. 2. Results of the alum-bronze material under three lighting environments. The top right shows the ground truth cascaded rendering (a sphere rendered with different point source directions) of the alum-bronze material. Column (a) shows the ground truth alum-bronze material rendered with one of the three lighting environments, column (b) shows a rendering of the estimated BRDF with the next ground truth lighting environment, column (c) shows the estimated illumination map, column (d) shows the ground truth illumination map, and column (e) shows a cascaded rendering of the recovered reflectance. The lighting environments used were Kitchen (1), Eucalyptus Grove (2), and the Uffizi Gallery (3). We achieve good estimates of reflectance and illumination, although the recovered illumination is missing high frequency details lost during image formation.

where $\mathbf{N}$ are the surface normals and σ_i^2 is the variance of the noise.

Assuming a linear camera, the irradiance $\mathbf{E}$ is computed as the reflectance radiance by integrating the incident irradiance modulated by the reflectance over the illumination map

$$\mathbf{E}(\mathbf{\Psi}, \mathbf{L}, \mathbf{N_x}) = \int \varrho(t(\omega_i, \omega_o); \mathbf{\Psi}) \mathbf{L}(\omega_i) \max(0, \mathbf{N_x} \cdot \omega_i) \, d\omega_i \,, \tag{11}$$

where t is a function which transforms ω_i and ω_o into the alternate BRDF parameterization variables θ_d and θ_h.

4.1 MAP Estimation

With the complete description of the posterior distribution, we are now ready to jointly estimate reflectance and illumination by computing the MAP estimate. In practice, we do this in an EM iteration by estimating the variance of the Gaussian likelihood in the expectation step and then in the maximization step jointly estimating the reflectance and illumination. The maximization step consists of alternatively minimizing the negative log posterior with respect to illumination and reflectance, i.e., fixing one to the current estimate while updating the other. When minimizing the negative log posterior with respect to the DSBRDF coefficients $\mathbf{\Psi}$ we use the Levenberg-Marquardt algorithm [19]. For our experiments we chose an illumination map of size 128×64 which we believe provides a good balance between computational speed and recovery detail. The large number of parameters, however, precludes the use of the standard Levenberg-Marquardt algorithm because of its high memory requirement. Instead, we minimize

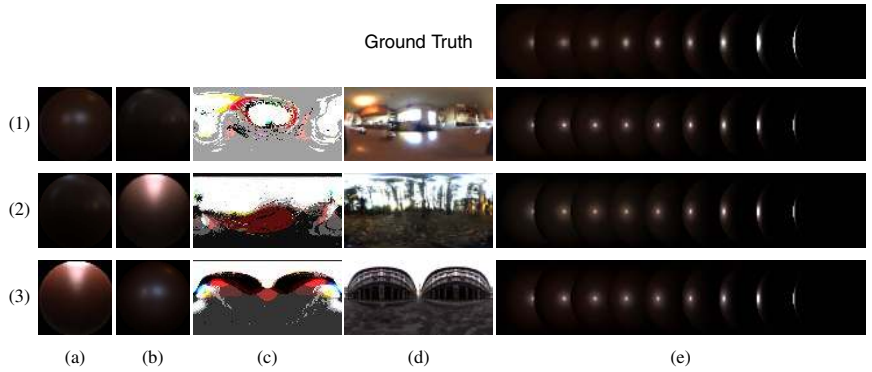

Fig. 3. Results for the cherry-235 material under three lighting environments. Results are presented as in Figure 2. As shown, we achieve accurate recovery of reflectance and illumination. Again, because of the low frequency reflectance of the cherry-235 material, there are even less high frequency details in the recovered illumination maps.

the error function with respect to the illumination map $\mathbf{L}$ using the limited-memory BFGS algorithm [20] which also allows us to apply a non-negativity constraint.

To initialize our algorithm, we set the entire illumination map $\mathbf{L}$ and the DSBRDF coefficients $\mathbf{\Psi}$ to zero. For the lighting environment this corresponds to pitch-black room, but for the reflectance it corresponds to the mean BRDF of the database used to compute the basis functions. This gives us a useful starting point for the BRDF, and we estimate lighting first.

4.2 Resolving the Color Ambiguity

One problem we have not addressed thus far is the color constancy problem between the reflectance and the illumination. As previously mentioned, if we are given an image of an object that appears red, we cannot know whether the red appearance is due to the reflectance or the illumination. This is made even more difficult by the fact that we are only dealing with a single object: if we observed that all the objects in a scene were red, we might assume that the redness is due to illumination. To handle the color ambiguity, we adopt the grey-world hypothesis [21] and assume that the lighting environment is, on average, uncolored. Consequently, we assume that the dominant coloration of the observed image is caused by reflectance.

We propose a simple way of exploiting this assumption. When performing joint estimation, we run our estimation algorithm twice: once constraining the illumination map to be greyscale, and again using the previous result as the initial input. During the first run this has the effect of forcing the BRDF to explain all the color variation in the input image. The colored BRDF is then used as an initial estimate for the second run of the algorithm. In all the lighting environments and materials we tried this approach and it successfully deduced the material color from the scene color.

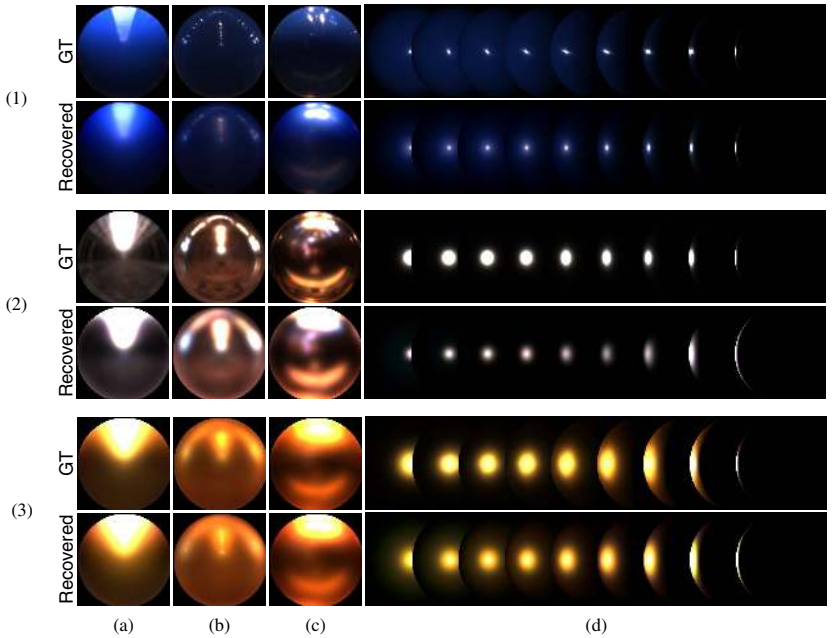

Fig. 4. Quality of reflectance estimates. Each subfigure demonstrates the reflectance estimates for the blue-acrylic (1), nickel (2), and gold-metallic-paint (3) materials. The top row of each subfigure shows the ground truth and the bottom row shows our estimates. Columns (a), (b), and (c) show renderings using the Uffizi Gallery, St. Peter's Basilica, and Grace Cathedral lighting environments, respectively. Column (d) is a cascaded rendering of the material with a series of point lights. The top left image of each subfigure was used as the input image. These results demonstrate the accuracy of our reflectance estimation.

5 Experimental Results

We apply our joint estimation approach to a number of synthetic and real scenes to evaluate its effectiveness. For the synthetic inputs, we render a representative selection of BRDFs from the MERL database [18] under a range of natural illuminations [22] to understand how performance varies. For real scenes, we compare directly to the work by Romeiro and Zickler [10] and show that we can explicitly estimate the illumination and both the reflectance and illumination in full color. In addition, we show results on our own data set.

5.1 Synthetic Results

Figures 2 and 3 show the results of our method for three different BRDFs from the MERL/MIT database [18] using three different lighting environments [22]. The results demonstrate that we achieve an accurate decomposition of lighting and reflectance when compared with ground truth. We note that, as we discussed, we are unable to recover a perfect lighting environment because of the bandpass effects of the BRDF. We still,

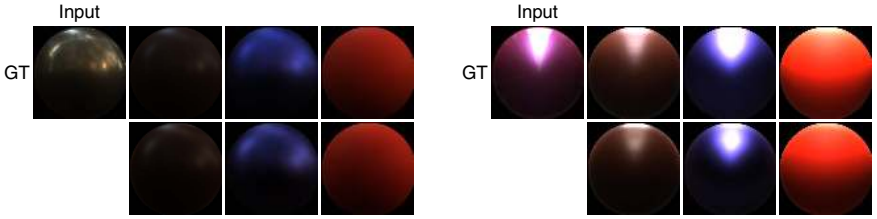

(a) Alum-bronze under Eucalyptus Grove **(b)** Violet-acrylic under Uffizi Gallery

Fig. 5. Predicting the appearance of materials with recovered illumination. We use the recovered illumination map to predict the appearance of materials with lower frequency reflectance. The top row shows the ground truth and the bottom row shows the predicted appearances. The input image for each subfigure is the top-left image. These results demonstrate the ability to accurately predict object appearance with the recovered illumination map.

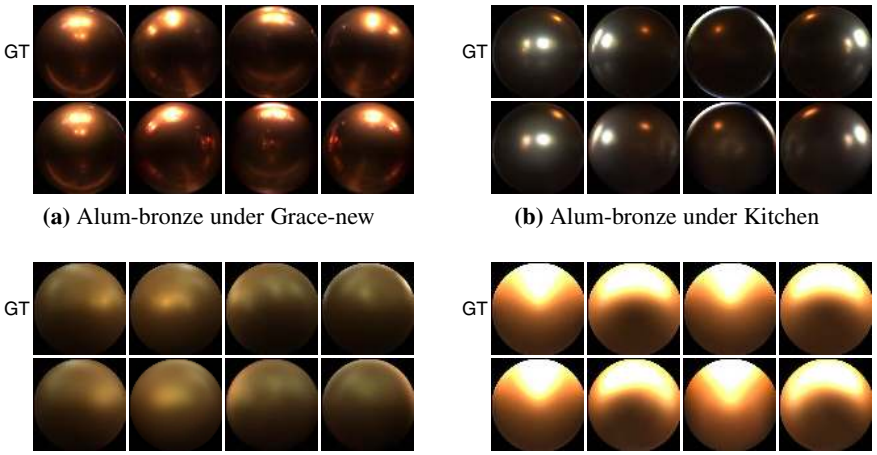

(a) Alum-bronze under Grace-new **(b)** Alum-bronze under Kitchen

(c) Gold-paint under Eucalyptus Grove **(d)** Gold-paint under Uffizi Gallery

Fig. 6. Predicting object appearance from different views. By using the recovered illumination map and reflectance estimate, we can show the object as it would appear from different viewing directions. The first row of each subfigure contains ground truth renderings for various BRDFs, using the measured BRDF values and true environment maps, and the second row contains the predicted appearance using the recovered reflectance and illumination from our method.

however, achieve a good and usable estimate that is only missing the high frequency information lost during image formation.

Figure 4 shows the results of three additional BRDFs under the Uffizi Gallery lighting environment. We demonstrate the quality of the reflectance estimates by rendering the estimated reflectance with the St. Peter Basilica and Grace Cathedral lighting environments, as well as a cascaded rendering of the BRDF under point lights. The results on blue-acrylic show that we correctly estimate the reflectance as a blue diffuse lobe plus a white specular lobe. This ability gives us an advantage over the work by

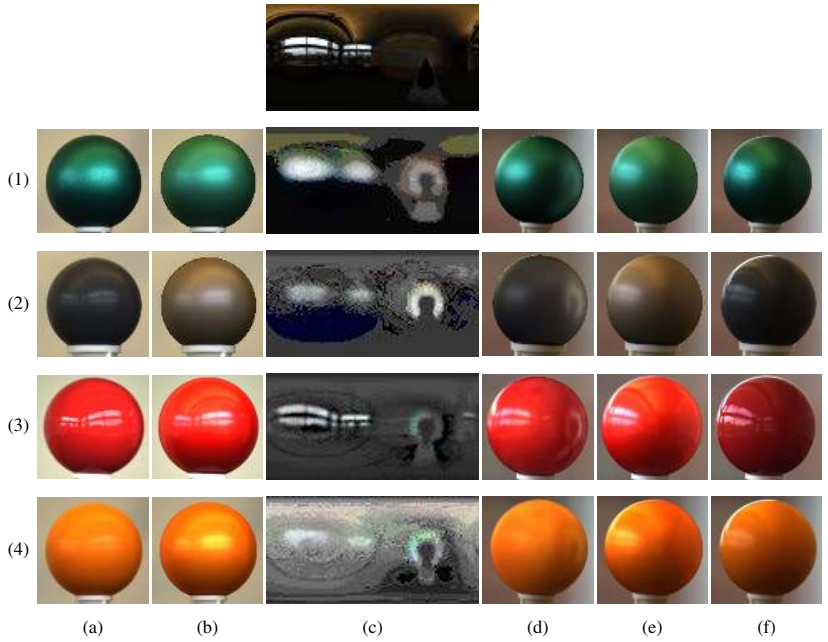

Fig. 7. Results on the data set from Romeiro and Zickler [10]. The top row is the ground truth illumination; rows 1–4 show the results on each material. Column (a) shows the input image, column (b) shows the recovered BRDF rendered with the ground truth illumination, column (c) shows the recovered illumination map, column (d) shows the recovered BRDF rendered with the recovered illumination map rotated to the alternate view, column (e) shows the recovered BRDF rendered with the ground truth alternate view illumination, and column (f) shows the ground truth alternate view. As we do not know the exact view point for the alternate view, there are slight differences in the viewpoints. The results demonstrate our method's ability to properly estimate a BRDF in full color. The method by Romeiro and Zickler (see figure 6 of [10]) is unable to capture the white highlights of the red sphere because they only estimate a monochrome BRDF.

Romeiro and Zickler [10] who only estimate a monochrome BRDF. This figure also demonstrates the method's ability to correctly estimate a diverse array of material reflectance (i.e., diffuse + specular, mirror-like, and metallic).

Figure 5 demonstrates how we may use the recovered illumination map to accurately predict the appearance of other materials. The results show that, even for materials that are not close to perfect mirror reflectance, we can use the recovered illumination to accurately compute the appearance of other materials, showing the ability of using an object of arbitrary material as a light probe.

Figure 6 shows one of our most important results: that we can accurately predict object appearance in the same scene from different viewing directions. The figure shows renderings of the ground truth BRDFs and illumination compared with our recovered reflectance and illumination from different viewing directions. The results show the ability to properly predict object appearance from any location in a scene. As mentioned, this ability could benefit object recognition and tracking algorithms.

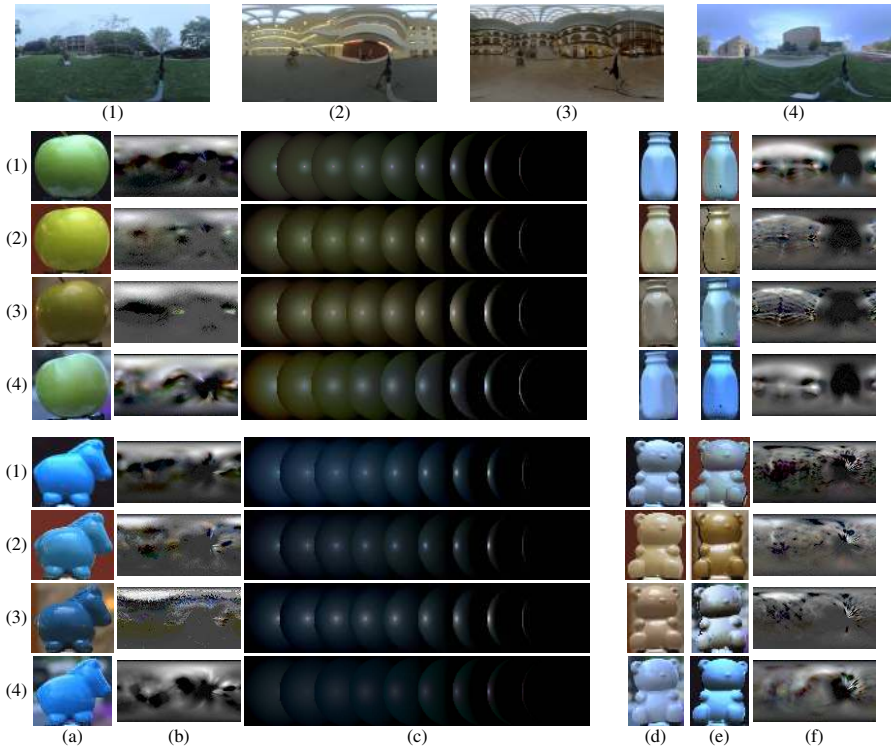

Fig. 8. Results on our new data set which includes four different objects under four different illumination environments (1–4). Columns (a), (b), and (c) show results for the apple and horse object in the four lighting environments. Column (a) shows the input image, column (b) shows the recovered illumination map, and column (c) shows a cascaded rendering of the recovered BRDF. Columns (d), (e), and (f) show the results for the milk bottle and bear object in the four lighting environments. Column (d) shows the input image, column (e) shows the recovered BRDF relit with the *next* ground truth illumination map, and column (f) shows the recovered illumination map. Our method recovers a good estimate of the band-limited lighting environment and BRDF despite some errors in the white balance of the camera which cause some small inaccuracies in the color of the recovered BRDF.

5.2 Real Results

Figure 7[1] shows results on the data set captured by Romeiro and Zickler [10]. The results show that recovered BRDF estimates closely match their ground truth (see the spheres rendered with recovered BRDF but with ground truth illumination from a different view in column (e)). They also demonstrate that we can accurately predict how the object would appear from a different viewpoint without any a priori knowledge of

[1] The ring on the right side of the recovered illumination is an artifact due to the matting at the occluding boundary. The light from the background is partially mixed into the boundary intensities resulting in a hallucinated illumination. This can be corrected with an alpha matte if necessary.

the true BRDF and lighting (compare column (d) with column (f)). In some cases (e.g., the green-metallic-paint) the predicted appearance using both the recovered BRDF and illumination looks closer to the ground truth than the one rendered with ground truth illumination. This is due to the fact that any over-estimation in the BRDF is compensated in the recovered illumination and thus when combined better match the true appearance. In addition, by estimating a full-color BRDF and illumination, we can recover and predict colors of the environment found in the highlights of the specular materials, which was not possible in the past [10].

Figure 8 shows the results of our method on a new database of high dynamic-range images of real objects with aligned ground-truth geometry taken under a variety of indoor and outdoor illumination environments that we introduce[2]. The results indicate that we are able to recover a good estimate of the BRDF of each example despite some error in hue caused by the color balance of the original photograph (i.e., the blue appearance of the outdoor images). The estimated illumination maps accurately capture the parts of the scene that contribute the most light, such as the sky and ceiling lights. This is true for all the objects except the milk bottle, whose surface normals do not span the entire hemisphere. Because of this, our method can only estimate detail in certain areas of the lighting map, leaving the rest relatively uniform in color.

These results demonstrate the method's ability to achieve accurate performance under a variety of lighting conditions and object types.

6 Conclusion

In this paper, we introduced a novel method to jointly estimate the reflectance and natural illumination from a single image of an object of known geometry. We showed how the space of illumination and reflectance can be constrained and incorporated as statistical and information theoretic priors into a Bayesian formulation. In particular, we introduced an entropy prior on the illumination to overcome the ambiguity between illumination and reflectance introduced by the image formation process. Our experimental results validate our choice of priors and formulation and demonstrate the effectiveness of the method on estimating both the reflectance and natural illumination from a single image in full color. In future work, we plan to extend the method to handle object surfaces consisting of multiple materials.

Acknowledgments. This work was supported in part by the Office of Naval Research grant N00014-11-1-0099, and the National Science Foundation awards IIS-0746717 and IIS-0964420. The authors also thank Kenji Hara for early discussions.

References

1. Ramamoorthi, R., Hanrahan, P.: A signal-processing framework for inverse rendering. In: Proc. of ACM SIGGRAPH, pp. 117–128 (2001)
2. Chandraker, M., Ramamoorthi, R.: What an image reveals about material reflectance. In: ICCV, pp. 1–8 (2011)

[2] Available online at `http://www.cs.drexel.edu/~kon/natgeom`

3. Lensch, H.P.A., Kautz, J., Goesele, M., Heidrich, W., Seidel, H.P.: Image-based Reconstruction of Spatial Appearance and Geometric Detail. ACM Trans. on Graphics 22(2), 234–257 (2003)
4. Zheng, Q., Chellappa, R.: Estimation of illuminant direction, albedo, and shape from shading. IEEE Trans. on Pattern Analysis and Machine Intelligence 13, 680–702 (1991)
5. Zickler, T., Ramamoorthi, R., Enrique, S., Belhumeur, P.N.: Reflectance Sharing: Predicting Appearance from A Sparse Set of Images of a Known Shape. IEEE Trans. on Pattern Analysis and Machine Intelligence 28(8), 1287–1302 (2006)
6. Hara, K., Nishino, K., Ikeuchi, K.: Mixture of Spherical Distributions for Single-View Relighting. IEEE Trans. on Pattern Analysis and Machine Intelligence 30(1), 25–35 (2008)
7. Hara, K., Nishino, K.: Variational Estimation of Inhomogeneous Specular Reflectance and Illumination from a Single View. Journal of Optical Society America, A 28(2), 136–146 (2011)
8. Marschner, S., Greenberg, D.: Inverse lighting for photography. In: IS&T/SID Fifth Color Imaging Conference, pp. 262–265. The Society for Imaging Science and Technology (1997)
9. Nishino, K., Ikeuchi, K., Zhang, Z.: Re-rendering from a sparse set of images. Technical Report DU-CS-05-12, Dept. of Computer Science, Drexel University (2005)
10. Romeiro, F., Zickler, T.: Blind Reflectometry. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010, Part I. LNCS, vol. 6311, pp. 45–58. Springer, Heidelberg (2010)
11. Nishino, K.: Directional Statistics BRDF Model. In: ICCV, pp. 476–483 (2009)
12. Nishino, K., Lombardi, S.: Directional Statistics-based Reflectance Model for Isotropic Bidirectional Reflectance Distribution Functions. Journal of Optical Society America, A 28(1), 8–18 (2011)
13. Huang, J., Mumford, D.: Statistics of natural images and models. In: CVPR, pp. 541–547 (1999)
14. Alldrin, N.G., Mallick, S.P., Kriegman, D.J.: Resolving the generalized bas-relief ambiguity by entropy minimization. In: CVPR, pp. 1–7 (June 2007)
15. Finlayson, G.D., Drew, M.S., Lu, C.: Entropy minimization for shadow removal. Int'l Journal of Computer Vision 85, 35–57 (2009)
16. Rusinkiewicz, S.: A New Change of Variables for Efficient BRDF Representation. In: Eurographics Workshop on Rendering, pp. 11–22 (1998)
17. Lombardi, S., Nishino, K.: Single image multimaterial estimation. In: CVPR, pp. 238–245 (2012)
18. Matusik, W., Pfister, H., Brand, M., McMillan, L.: A Data-Driven Reflectance Nodel. ACM Trans. on Graphics 22(3), 759–769 (2003)
19. Marquardt, D.W.: An algorithm for least-squares estimation of nonlinear parameters. SIAM Journal on Applied Mathematics 11(2), 431–441 (1963)
20. Byrd, R.H., Lu, P., Nocedal, J., Zhu, C.: A Limited Memory Algorithm for Bound Constrained Optimization. SIAM Journal on Scientific Computing 16, 1190–1208 (1995)
21. Buchsbaum, G.: A spatial processor model for object colour perception. Journal of the Franklin Institute 310(1), 1–26 (1980)
22. Debevec, P.: Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In: Proc. of ACM SIGGRAPH, pp. 189–198 (1998)