

Estimating Intrinsic Images from Image Sequences with Biased Illumination

Yasuyuki Matsushita¹, Stephen Lin¹, Sing Bing Kang², and
Heung-Yeung Shum¹

¹ Microsoft Research Asia
3F, Beijing Sigma Center, No.49, Zhichun Road Haidian District,
Beijing 100080, China

² Microsoft Research
One Microsoft Way, Redmond, Washington 98052-6399, U.S.A.
{yasumat, stevelin, sbkang, hshum}@microsoft.com

Abstract. We present a method for estimating intrinsic images from a fixed-viewpoint image sequence captured under changing illumination directions. Previous work on this problem reduces the influence of shadows on reflectance images, but does not address shading effects which can significantly degrade reflectance image estimation under the typically biased sampling of illumination directions. In this paper, we describe how biased illumination sampling leads to biased estimates of reflectance image derivatives. To avoid the effects of illumination bias, we propose a solution that explicitly models spatial and temporal constraints over the image sequence. With this constraint network, our technique minimizes a regularization function that takes advantage of the biased image derivatives to yield reflectance images less influenced by shading.

1 Introduction

Variation in appearance caused by illumination changes has been a challenging problem for many computer vision algorithms. For example, face recognition is complicated by the wide range of shadow and shading configurations a single face can exhibit, and image segmentation processes can be misled by the presence of shadows and shading as well. Since image intensity arises from a product of reflectance and illumination, one approach for dealing with variable lighting is to decompose an image into a reflectance component and an illumination component [7], also known as intrinsic images [1]. The reflectance image, free of illumination effects, can then be processed without consideration of shadows and shading.

Decomposition of an image into intrinsic images, however, is an underconstrained problem, so previous approaches in this area introduced additional constraints to make the problem tractable. In [6], it is assumed that the illumination component is spatially smooth while the reflectance component exhibits sharp changes, such that low-pass filtering of the input image yields the illumination image. Similarly, [3] assumes smoothness of illumination and piecewise constant reflectance, so that removing large derivatives in the input image results in the illumination image. In addition to illumination smoothness, Kimmel *et al.* [5] include constraints that the reflectance is smooth and the illumination image is close to the input image.

Instead of relying on smoothness constraints, Tappen *et al.* [10] proposed a learning-based approach to separate reflectance edges and illumination edges in a derivative image. Although this method successfully separates reflectance and shading for a given illumination direction used in training, it is difficult to create such a prior to classify edges under arbitrary lighting. Another edge-based method was proposed by Finlayson *et al.* [11] that suppresses color derivatives along the illumination temperature direction to derive a shadow-free image of the scene. In addition to shadow edges, this approach may remove texture edges that also have a color derivative in the illumination temperature direction.

Rather than using only a single input image, Weiss [9] deals with the simpler scenario of having an image sequence captured from a fixed viewpoint with changing illumination conditions. This method employs a ML estimation framework based on a prior that illumination images yield a Laplacian distribution of derivatives between adjacent pixels. Experimental results demonstrate that this technique efficiently and robustly removes cast shadows from reflectance images. Shading on non-planar surfaces, however, can significantly degrade ML estimation of intrinsic images by altering the distribution of derivatives, especially in the typical case of a biased illumination distribution that is not centered around the surface normals of the adjacent pixel pair. More recently, Matsushita *et al.* [12] extended Weiss’s method to handle the scenes where the Lambertian assumption does not hold. Using the reflectance image estimated by ML estimation as a scene texture image, their method derives time-varying reflectance images instead of assuming a single reflectance image.

In our proposed method, we also take as input an image sequence and analyze the derivative distributions. Because of the effects of illumination bias on the derivative distributions, we present an alternative method for processing image derivatives, based on explicit modeling of spatial and temporal constraints over the image sequence. With this constraint network, a reflectance image and a set of illumination images are estimated by minimizing a function based on smoothness of illumination and reflectance. Although the derivative distributions are unsuitable for ML estimation, our technique nevertheless takes advantage of derivative distribution information to spatially vary the weight of the smoothness constraints in a manner unlike previous regularization-based methods.

The goal of this work is closely related to that of photometric stereo with unknown light sources and spatially varying albedo. One strong assumption in most uncalibrated photometric stereo approaches [19–21] is that there are no cast shadows. However, it is clear that this assumption does not hold in many situations for real world scenes. Yuille *et al.* [22] have proposed a method to handle cast shadows using robust statistics; however, one drawback of the method is that it assumes a single point source in each image. Photometric stereo yields accurate results, but generally it is necessary to assume limiting conditions. While the photometric stereo framework relies on the structural smoothness, our method relies more on the smoothness of reflectance and illumination images. In the context of photometric stereo, Wolff and Angelopoulou [4] acquired multiple stereo pairs of images of the same scene under different illumination conditions. With two stereo pairs they ob-

tain a stereo pair of photometric ratio images, in which the albedo term is removed in order to extend geometric stereo reconstruction to smooth featureless surfaces.

The remainder of the paper is organized as follows. Sec. 2 details the problem of illumination bias and the resulting effects of shading on derivative distributions. In Sec. 3, we describe the constraints on the energy minimization process and the influence of the derivative distribution. Our algorithm is presented in Sec. 4, followed by experimental results in Sec. 5 and a conclusion in Sec. 6.

2 Effect of Illumination Bias

Before describing the effect of illumination bias on derivative distributions, we begin with a brief review of intrinsic images and the ML estimation technique.

Under the Lambertian assumption, as expressed in the following equation, an input image I arises from a product of two intrinsic images: the reflectance image ρ and the illumination image L . Since the viewpoint of the image sequence is fixed, ρ does not vary with time t . The illumination is comprised by an ambient term α and a direct illumination term L_D , which is the product of the illumination intensity E , a binary shadow presence function g and the inner product of surface normal \mathbf{n} and illumination direction \mathbf{l} :

$$\begin{aligned} I(x, y, t) &= \rho(x, y)L(x, y, t) \\ &= \rho(x, y)\{L_D(x, y, t) + \alpha(x, y, t)\} \\ &= \rho(x, y)\{E(t)g(x, y, t)(\mathbf{n}(x, y) \cdot \mathbf{l}(t)) + \alpha(x, y, t)\} \\ &= \rho(x, y)E(t)\{g(x, y, t)(\mathbf{n}(x, y) \cdot \mathbf{l}(t)) + \alpha'(x, y, t)\} \end{aligned} \quad (1)$$

where $\mathbf{n} \cdot \mathbf{l}$ is always non-negative, and α' indicates the ambient term normalized by the illumination intensity E . In the ML estimation framework of [9], n derivative filters f_n are first applied to the logarithms of images $I(t)$. For each filter, a filtered reflectance image ρ_n is then computed as the median value in time of $f_n \star \log I$, where \star represents convolution:

$$\log \hat{\rho}_n(x, y) = \text{median}_t \{f_n \star \log I(x, y, t)\}. \quad (2)$$

The filtered illumination images $\log L_n(x, y, t)$ are then computed using the estimated filtered reflectance images $\log \hat{\rho}_n$ according to

$$\log \hat{L}_n(x, y, t) = f_n \star \log I(x, y, t) - \log \hat{\rho}_n(x, y). \quad (3)$$

Finally, a reflectance image ρ and illumination images L are recovered from the filtered reflectance images ρ_n and illumination images $L_n(t)$ through the following deconvolution process,

$$(\log \hat{\rho}, \log \hat{L}) = h \star \left(\sum_n f_n^r \star (\log \hat{\rho}_n, \log \hat{L}_n) \right) \quad (4)$$

where f_n^r is the reversed filter of f_n , and h is the filter which satisfies the following equation:

$$h \star \left(\sum_n f_n^r \star f_n \right) = \delta. \quad (5)$$

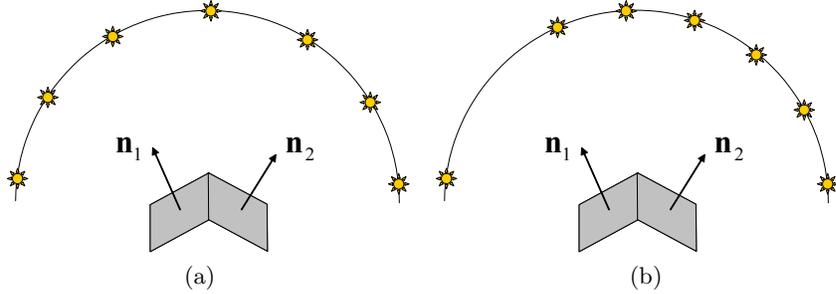


Fig. 1. Illumination conditions. (a) Uniform illumination, (b) biased illumination.

From (2), it can be shown that for two adjacent pixels with intensities $I_1(t)$ and $I_2(t)$,

$$\hat{\rho}_n = \text{median}_t \frac{I_1(t)}{I_2(t)} = \text{median}_t \frac{\rho_1}{\rho_2} \cdot \frac{E(t)\{g_1 \cdot (\mathbf{n}_1 \cdot \mathbf{l}(t)) + \alpha'_1\}}{E(t)\{g_2 \cdot (\mathbf{n}_2 \cdot \mathbf{l}(t)) + \alpha'_2\}}. \quad (6)$$

We assume that α' is constant over an image, i.e., $\alpha'_1(t) = \alpha'_2(t)$. Cast shadows affect this equation only when $g_1 \neq g_2$. Since this instance seldom occurs, cast shadows do not affect the median derivative values used in ML estimation. It can furthermore be seen that when $\mathbf{n}_1 = \mathbf{n}_2$, shading does not affect ML estimation since $\mathbf{n}_1 \cdot \mathbf{l} = \mathbf{n}_2 \cdot \mathbf{l}$ and consequently $\hat{\rho}_n = \rho_1/\rho_2$.

When a pair of pixels have different surface normals, ML estimation can also deal with shading in cases of *unbiased* illumination samples. For a pair of adjacent pixels with surface normals \mathbf{n}_1 and \mathbf{n}_2 , the set Ω_l of illumination samples $\mathbf{l}(t)$ are unbiased only under the following condition:

$$\text{median}_{\mathbf{l}(t) \in \Omega_l} \{\mathbf{n}_1 \cdot \mathbf{l}(t) - \mathbf{n}_2 \cdot \mathbf{l}(t)\} = 0. \quad (7)$$

In other words, the illumination is unbiased for a given pair of pixels when the illumination image value $L(x, y)$ of both pixels is the same for the median derivative value. Otherwise, the illumination distribution is biased. Figure 1 shows an illustration of unbiased illumination and biased illumination for a given pair of pixels. With unbiased illumination as given in (7), it can be seen that (6) results in the correct value ρ_1/ρ_2 .

When a pair of adjacent pixels have different surface normals, illumination bias will cause the ML estimation to be incorrect, because $\mathbf{n}_1 \cdot \mathbf{l} \neq \mathbf{n}_2 \cdot \mathbf{l}$ for the median derivative value. In this case, the illumination ratio in (6) does not equal to one, and consequently $\hat{\rho}_n \neq \rho_1/\rho_2$. This can be expected since different shading is present in the two pixels for every observation.

The case of different surface normals with illumination bias is a significant one, because for a pair of adjacent non-planar pixels, unbiased illumination is rare. So for most pairs of non-planar pixels, ML estimation fails to compute the correct reflectance ratio and the estimated reflectance image contains shading. Figure 2 shows a typical result of ML estimation applied to a synthetic scene with non-planar surfaces. A ball on a plane is lit from various directions as exemplified in

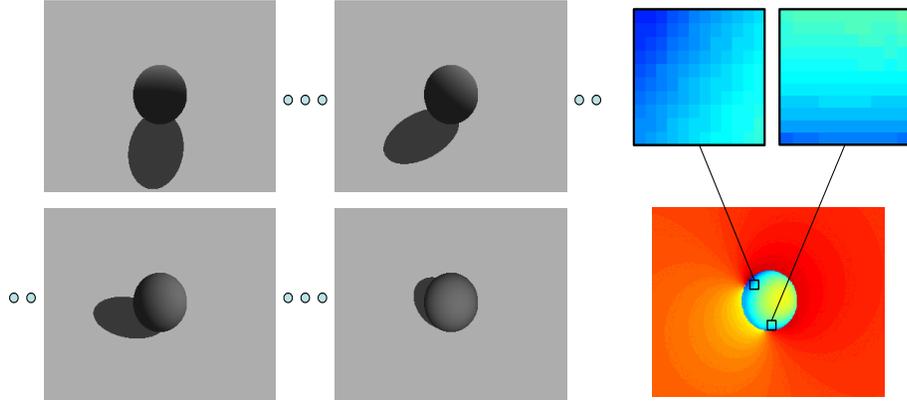


Fig. 2. Shading effect remains on reflectance estimate with ML estimation. Left: an input image sequence, Right: the estimated reflectance image with ML estimation.

the input images on the left side of Figure 2. Although the illumination samples are unbiased for some pairs of pixels, they are biased for most pairs of adjacent pixels. As a result, shading remains in the estimated reflectance image as shown on the right side of the figure.

3 Solution Constraints

Since ML estimation is generally affected by shading, we propose an alternative solution method based on the constraints described in this section. Let us denote i, j as labels for illumination directions, p, q for adjacent pixels, and N, M for the number of observed illumination conditions and the number of pixels in an image, respectively.

From a sequence of images, we can derive spatial constraints between adjacent pixels (inter-pixel) and temporal constraints between corresponding pixels under different light directions (inter-frame). Moreover, we employ smoothness constraints to make the problem tractable.

[Inter-frame constraint]

Assuming that the scene is composed of Lambertian surfaces, the reflectance value at each point ρ is constant with respect to time. We can thereby derive a temporal constraint as follows:

$$\frac{I_p(t_i)}{I_p(t_j)} = \frac{L_p(t_i)}{L_p(t_j)}, \quad 0 \leq i, j < N; i \neq j. \quad (8)$$

This does not determine the absolute values of L s; however, it fixes the ratios among L_p s.

[Inter-pixel constraint]

Letting ω_p be a set of pixels that are neighbours of p ,

$$\frac{I_p(t_i)}{I_q(t_i)} = \frac{\rho_p}{\rho_q} \cdot \frac{L_p(t_i)}{L_q(t_i)}, \quad 0 \leq i < N; q \in \omega_p. \quad (9)$$

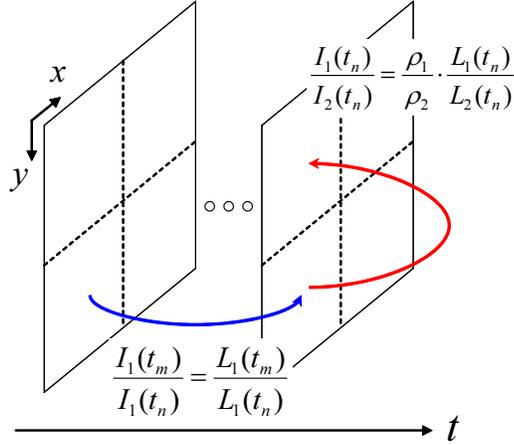


Fig. 3. Inter-frame and inter-pixel constraints. A set of constraints composes a constraint network.

This constraint can possibly be applied to non-neighboring pixels, however, we restrict this to be applied only to neighboring pixels because we use the flatness constraint and smooth reflectance constraint in the energy minimization step.

These constraints can be derived from (1), and they compose a 3-D constrained network about L and a 2-D constrained network about ρ as illustrated in Fig. 3. We use them as hard constraints and force ρ and L to always satisfy the following equation:

$$\sum_{p,i,j;i \neq j} \left(\frac{I_p(t_i)}{I_p(t_j)} - \frac{L_p(t_i)}{L_p(t_j)} \right)^2 + \sum_{p,q,i;q \in \omega_p} \left(\frac{I_p(t_i)}{I_q(t_i)} - \frac{\rho_p}{\rho_q} \cdot \frac{L_p(t_i)}{L_q(t_i)} \right)^2 = 0. \quad (10)$$

[Smoothness constraints]

In addition to the above constraints, our technique favors smoothness over both ρ and L . Smoothness is a generic assumption underlying a wide range of physical phenomena, because it characterizes coherence and homogeneity. Based on the fact that retinal images tend to be smooth according to natural image statistics [16], we assume that both ρ and L are smooth as well. By formulating these two assumptions into an energy function, we derive intrinsic images by an energy minimization scheme.

The choice of energy function $E(\rho, L)$ is a critical issue. Various kinds of energy functions that measure the smoothness of observed data have been proposed. For example, in regularization-based methods [13, 14], the energy minimization function makes the observed data smooth everywhere. This generally yields poor results at object boundaries since discontinuities should be preserved. One discontinuity-preserving function is Geman and Geman's MRF-based function [15].

Although our method assumes smoothness of L , this condition clearly does not hold when the surface normals of adjacent pixels differ. In such instances, the smooth

L constraint should be weakened. To estimate the amount of difference between neighboring surface normals, we use the information present in the derivative distribution.

If a pair of pixels lie on a flat surface, the values of $I_p(t)/I_q(t)$ are almost always equal to 1 except when only one of the pixels is shadowed, as discussed in [9]. We use this strong statistic and define an error function based on the hypothesis that I_p and I_q share the same surface normal:

$$e_{pq}(t_i) = \left| \arctan \left\{ \text{median}_t \left(\frac{I_p}{I_q} \right) \right\} - \arctan \left\{ \frac{I_p}{I_q} \right\} \right| \quad (11)$$

In (11), $\text{median}_t(I_p/I_q)$ corresponds to ML estimation, which gives the ratio of reflectance if p and q are co-planar. We evaluate the angle between the ratio of reflectance and the ratio of observed intensity to determine if the observation supports the flatness hypothesis. To determine the amount of support for the flatness hypothesis, a threshold ϵ is used:

$$\xi_{pq}(t_i) = \begin{cases} 1 & (e_{pq}(t_i) < \epsilon : \text{accept}) \\ 0 & (e_{pq}(t_i) \geq \epsilon : \text{reject}) \end{cases} \quad (12)$$

Finally, we compute the ratio of the number of acceptances ξ to the number of total observations N to test the surface-flatness hypothesis. When the surface flatness f is high, it is likely that L is smooth. On the other hand, when f is not high, the smoothness assumption for L should be weakened. We define the flatness f as the square of the acceptance ratio:

$$f_{pq} = \left(\frac{\sum_i \xi_{pq}(t_i)}{N} \right)^2. \quad (13)$$

Using the surface flatness evaluated by (13), we define an energy function E_Ω to minimize:

$$\begin{aligned} E_\Omega &= \sum_p E_p(\Delta\rho_p, \Delta L_p(t)) \\ &= \sum_p \sum_{q \in \omega_p} \{ (\rho_p - \rho_q)^2 + \lambda f_{pq}(t_i) (L_p(t_i) - L_q(t_i))^2 \} \end{aligned} \quad (14)$$

where λ is a coefficient that balances the smoothness of ρ and L .

Equation (14) always converges to a unique solution because E_Ω is convex with respect to ρ_p and L_p . This is confirmed by taking E_p 's Hessian H_p

$$H_p = \begin{bmatrix} \partial^2 E_p / \partial \rho^2 & \partial^2 E_p / \partial \rho \partial L \\ \partial^2 E_p / \partial L \partial \rho & \partial^2 E_p / \partial L^2 \end{bmatrix} = 2 \begin{bmatrix} \sum_{q \in \omega_p} 1 & 0 \\ 0 & \lambda \sum_{q \in \omega_p} f_{pq} \end{bmatrix}, \quad (15)$$

the leading principal minors of which are $2 \sum_{q \in \omega_p} 1 > 0$, $2\lambda \sum_{q \in \omega_p} f_{pq} > 0$ where $\lambda > 0$, $f_{pq} > 0$, so that the function E_p is strictly convex. Since the sum of convex functions is convex, E_Ω is also convex because $E_\Omega = \sum_p E_p$.

4 Algorithm

With the constraints described in the preceding section, our algorithm can proceed as follows.

[Step 1 : Initialization] Initialize ρ and L .

[Step 2 : Hard constraints] Apply the inter-frame and inter-pixel constraints expressed in (10). Since it is difficult to minimize the two terms in (10) simultaneously, we employ an iterative approach for minimization.

1. Inter-frame constraint. Update $L_p(t_i)$.

$$L_p(t_i) \leftarrow \sum_{j \neq i} \left(\frac{I_p(t_i)}{I_p(t_j)} L_p(t_j) \right) / (N - 1) \quad (16)$$

2. Inter-pixel constraint. Update $L_p(t_i)$ and ρ_p with ratio error β . Letting M_{ω_p} be the number of p 's neighboring pixels,

$$\beta_p(t_i) = \left(\sum_{q \in \omega_p} \frac{I_p(t_i)}{I_q(t_i)} \cdot \frac{\rho_q L_q(t_i)}{\rho_p L_p(t_i)} \right) / M_{\omega_p}. \quad (17)$$

Since the error ratio $\beta_p(t_i)$ can be caused by some unknown combination of ρ and L , we distribute the error ratio equally to both ρ and L in (18) and (20), respectively.

$$L_p(t_i) \leftarrow \sqrt{\beta_p(t_i)} L_p(t_i), \quad (18)$$

$$\beta_p = \left(\sum_i \beta_p(t_i) \right) / N, \quad (19)$$

$$\rho_p \leftarrow \sqrt{\beta_p} \rho_p. \quad (20)$$

3. Return to 1. unless Equation (10) is satisfied.

[Step 3 : Energy minimization]

Evaluate the energy function (14), and find ρ and L that lower the total energy. If the total energy can still decrease, update ρ and L , then go back to Step 2. Otherwise, we stop the iteration. By fixing ρ_q and L_q in (14), the energy minimization is performed for each E_p using the conjugate gradient method. The conjugate gradient method is an iterative method for solving linear systems of equations which converge faster than the steepest descent method. For further details of the algorithm, readers may refer to a well presented review [18].

5 Experimental Results

To evaluate our method, we carried out experiments over one synthetic image sequence and three real world image sequences. In these experiments, we used 11 different lighting conditions, and set $\lambda = 0.4$. We used $\epsilon = 0.02$ which is empirically obtained in Equation (12) for all experiments. The determination of ϵ has dependency on minimum signal-to-noise ratio. Starting with constant initial values, ρ and L are iteratively updated. There is no restriction about the initial values, however, we used flat images for their initial values because of the smoothness assumption.

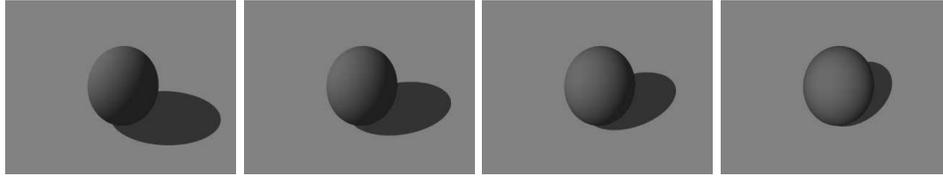


Fig. 4. Input image samples from synthetic scene. Illumination samples are chosen to be biased.

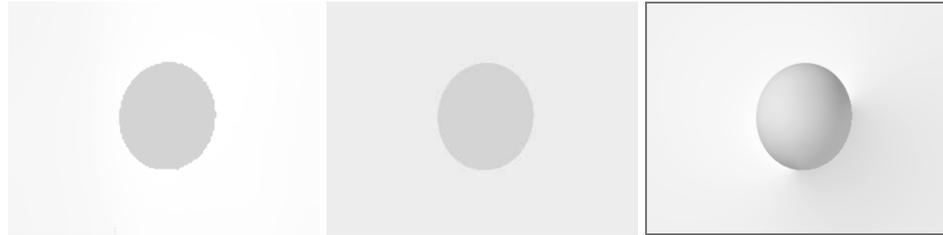


Fig. 5. Estimated reflectance images. Left : our method, Center : the ground truth, Right : ML estimation.

5.1 Synthetic scene

For a synthetic scene, we prepared a Lambertian scene with a sphere and a plane as shown in Figure 4. The illumination samples are biased in most cases, since most of them lie on the left-hand side of the images. Figure 5 shows the result of our method, the ground truth, and the result of ML estimation from left to right. Due to the scaling ambiguity of the reflectance images, we adjusted the scaling of each reflectance image for better comparison and display. As we can clearly see in the figure, our method can successfully derive shading-free a reflectance image which is close to the ground truth.

5.2 Real world scenes

For real world scenes, we captured two image sequences of toy scenes and used Yale Face Database B [17]. Figure 6 and Figure 8 show the result of reflectance estimation from Lambertian scenes. In both figures, the left image shows the result of our method, while the right image shows the result of ML estimation. As we can see clearly in them, our method handles shading more correctly and shading effect is much reduced in our reflectance estimates. Figure 7 and 9 show the estimated illumination images and corresponding input images. In illumination images, reflectance edges such as texture edges are well removed.

On the other hand, Figure 10 shows a negative result, especially on the hair. Since the human hair shows high specularity, and it is hard to model it as Lambertian. This non-Lambertian property affects to our method and turns those area into white. This is because our method is based on Lambertian model, and it implies that our method does not handle specular reflections well. However, as for the

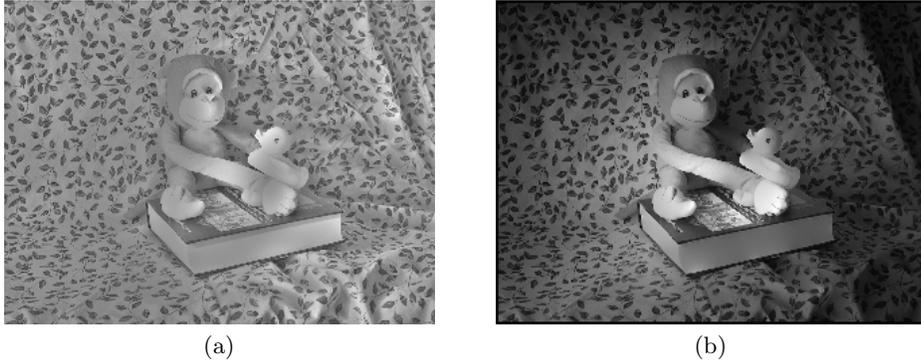


Fig. 6. Toy scene 1. Estimated reflectance images. (a) Our method, (b) ML estimation.



Fig. 7. (a) Estimated illumination image, and (b) the corresponding input image.

non-specular part such as the human face, shading effect is much reduced by our method.

6 Conclusion

We have presented a method that robustly estimates intrinsic images by energy minimization. Unlike previous methods, our method is not affected by illumination bias which generally exists. In our framework, we explicitly modeled spatial and temporal constraints over the image sequence to form a constraint network. Using this as a hard constraint, we minimized an energy function defined from the assumptions that reflectance and illumination are smooth. By weighting these smoothness constraints according to a surface flatness measure estimated from derivative distributions, we estimated intrinsic images with improved handling of shading. Evaluation with both synthetic and real world image sequences shows that our method can robustly estimate shading-free reflectance image and illumination images. Some of the next steps of our research will include the acceleration of energy minimization part and extension of our model to correctly handle specularity.

References

1. H.G. Barrow and J.M. Tenenbaum: Recovering intrinsic scene characteristics from images. In *A. Hanson and E. Riseman, editors, Computer Vision Systems. Academic*



Fig. 8. Toy scene 2. Estimated reflectance images. (a) Our method, (b) ML estimation.



Fig. 9. (a) Estimated illumination image, and (b) the corresponding original image.

Press, 3–26, 1978.

2. E.H. Adelson and A.P. Pentland: The Perception of Shading and Reflectance. In *D. Knill and W. Richards (eds.), Perception as Bayesian Inference*, 409–423, 1996.
3. A. Blake: Boundary conditions of lightness computation in mondrian world. In *Computer Vision, Graphics and Image Processing*, 32, 314–327, 1985.
4. L. B. Wolff and E. Angelopoulou: 3-d stereo using photometric ratios. In *Proc. of the Third European Conference on Computer Vision (ECCV)*, pp 247–258, 1994.
5. R. Kimmel, M. Elad, D. Shaked, R. Keshet and I. Sobel: A Variational Framework for Retinex. In *International Journal of Computer Vision*, 52(1), 7–23, 2003.
6. E.H. Land: An alternative technique for the computation of the designer in the Retinex theory of color vision. In *Proc. Nat. Acad. Sci.*, 83, 3078–3080, Dec. 1986.
7. E.H. Land: The Retinex theory of color vision. In *Scientific American*, 237(G), No. 6, 108–128, Dec. 1977.
8. E.H. Land, and J.J. McCann: Lightness and retinex theory. In *Journal of the Optical Society of America*, 61(1), 1–11, 1971.
9. Y. Weiss: Deriving intrinsic images from image sequences. In *Proc. of 9th IEEE Int'l Conf. on Computer Vision*, 68–75, Jul., 2001.
10. M.F. Tappen, W.T. Freeman, and E.H. Adelson: Recovering Intrinsic Images from a Single Image. In *Advances in Neural Information Processing Systems 15 (NIPS)*, MIT Press, 2002.
11. G.D. Finlayson, S.D. Hordley, and M.S. Drew: Removing Shadows from Images. In *Proc. of European Conf. on Computer Vision Vol.4*, 823–836, 2002.

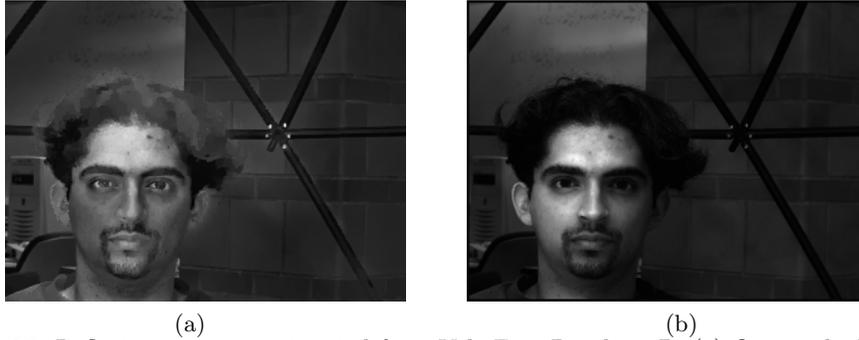


Fig. 10. Reflectance images estimated from Yale Face Database B. (a) Our method, (b) ML estimation. The result shows the limitation of our method, i.e. high specularity affects the result.

12. Y. Matsushita, K. Nishino, K. Ikeuchi and M. Sakauchi : Illumination Normalization with Time-dependent Intrinsic Images for Video Surveillance. In *Conf. on Computer Vision and Pattern Recognition (CVPR)*, Vol.1, pp. 3–10, 2003.
13. A. Blake: Comparison of the efficiency of deterministic and stochastic algorithms for visual reconstruction. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(1):2–12, 1989.
14. B.K.P. Horn and B. Schunk: Determining optical flow. In *Artificial Intelligence*, 17:185–203, 1981.
15. S. Geman and D. Geman: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
16. B.A. Olshausen and D.J. Field: Emergence of simplecell receptive field properties by learning a sparse code for natural images. In *Nature*, 381:607-608, 1996.
17. A.S. Georghiades, P.N. Belhumeur, and D.J. Kriegman: From few to many: Generative models for recognition under variable pose and illumination. In *IEEE Int. Conf. on Automatic Face and Gesture Recognition*, 277-284, 2000.
18. J.R. Shewchuck: An introduction to the conjugate gradient method without agonizing pain. *Tech. Rep. CMU-CS-94-125*, Carnegie Mellon University, 1994.
19. H. Hayakawa: Photometric stereo under a light-source with arbitrary motion. In *Journal of Optical Society of America A.*, 11(11):3079–3089, 1994.
20. R. Basri and D. Jacobs: Photometric stereo with general, unknown lighting. In *Proc. of Computer Vision and Pattern Recognition(CVPR)*, Vol.2, pp. 374–381, 2001.
21. A.S. Georghiades, P.N. Belhumeur and D.J. Kriegman: Illumination-Based Image Synthesis: Creating Novel Images of Human Faces Under Differing Pose and Lighting. In *Proc. Workshop on Multi-View Modeling and Analysis of Visual Scenes*, pp. 47–54, 1999.
22. A.L. Yuille, D. Snow, R. Epstein, P. Belhumeur: Determining Generative Models for Objects Under Varying Illumination: Shape and Albedo from Multiple Images Using SVD and Integrability. In *International Journal on Computer Vision.*, 35(3), pp 203–222. 1999.