Inverse Lighting from Cast Shadows under Unknown Radiometric Response Function

Takuto Nakashima, Ryo Matsuoka[0000−0003−4774−1183], and Takahiro Okabe[0000−0002−2183−7112]

Department of Artificial Intelligence
Kyushu Institute of Technology
okabe@ai.kyutech.ac.jp

Abstract. Inverse lighting is a technique for estimating the illumination distribution of a scene from a single image. Conventionally, inverse lighting assumes either a linear radiometric response function or a convex scene without cast shadows. Unfortunately, however, consumer cameras usually have unknown and nonlinear radiometric response functions, and then the existing methods do not work well for images such as Internet photos taken by using those cameras. Moreover, it is known that the high-frequency components of an illumination distribution cannot be recovered from diffuse reflection components without cast shadows. In this paper, we propose a method for jointly estimating both an illumination distribution and a radiometric response function from cast shadows. Specifically, our proposed method represents the illumination distribution and the response function by using Haar wavelets with the sparseness constraint and polynomials respectively, and then estimates their coefficients. We conducted a number of experiments by using both synthetic and real images, and confirmed that our method works better than the existing methods. In addition, we showed that masking pixels near edges makes the joint estimation robust for real images with approximate geometry.

Keywords: Inverse Lighting, Cast Shadows, Radiometric Response Function, Haar Wavelets, Spherical Harmonics

1 INTRODUCTION

The appearance of an object depends on the shape and reflectance of the object as well as the illumination distribution of a scene. Inverse lighting [6] is a technique for estimating the illumination distribution of a scene from a single input image with known shape and reflectance. Illumination recovery is important for applications to Augmented Reality (AR) and Mixed Reality (MR); it enables us to superimpose photometrically consistent synthetic objects onto real scenes and real images [1, 11]. Inverse lighting is important also as a building block of multi-view inverse rendering [15, 13, 16, 4]; it jointly estimates the detailed shape, reflectance, and illumination from multiple images such as Internet photos.

Conventionally, inverse lighting assumes either a linear radiometric response function or a convex scene. The former means that a pixel value is proportional to a radiance value. Unfortunately, however, consumer cameras usually have unknown and nonlinear radiometric response functions [3]. Therefore, the existing methods assuming a linear response function [6, 10, 12, 9] do not work well for images, in particular Internet photos, taken by using those cameras with unknown and nonlinear radiometric response functions.

The latter means that there are no shadows cast by one object onto another in the scene. It is known that the high-frequency components of an illumination distribution cannot be recovered from diffuse reflection components [10], but can be recovered from cast shadows [9]. Ohta and Okabe [8] propose a method for jointly estimating both the illumination distribution of a scene and the radiometric response function of a camera from diffuse reflection components, and therefore their method can recover only the low-frequency components of an illumination distribution. In addition, they tested their method only on synthetic images with accurate geometry, and then its effectiveness on real images with approximate geometry is questionable.

To cope with those limitations, we propose a method for jointly estimating both the illumination distribution of a scene and the radiometric response function of a camera from cast shadows. Specifically, our proposed method represents the illumination distribution and the response function by using Haar wavelets with the sparseness constraint and polynomials respectively, and then estimates their coefficients. We conducted a number of experiments using both synthetic and real images, and confirmed that our method works better than the existing methods. In addition, we empirically show that masking the pixels near the edges in an input image makes the joint estimation robust for real images with approximate geometry.

The main contributions of this study are twofold. First, we propose a method for jointly estimating both an illumination distribution and a radiometric response function from cast shadows; it enables us to estimate the high-frequency components of the illumination distribution from a single input image captured by using a consumer camera with an unknown and nonlinear response function. Second, through the experiments using both synthetic and real images, we experimentally confirmed that our method works better than the existing methods. Specifically, we showed the effectiveness of the joint estimation and the use of Haar wavelets with the sparseness constraint. In addition, we showed that the joint estimation from real images with approximate geometry becomes robust when the pixels near the edges in an input image are removed from the estimation.

2 PROPOSED METHOD

In this section, we propose a method for jointly estimating both the illumination distribution of a scene and the radiometric response function of a camera from a single input image. Our proposed method assumes that the shape and reflectance

of a scene are known, and the scene is illuminated by distant light sources in a similar manner to the conventional inverse lighting. We discuss how to make the joint estimation robust for real images with approximate geometry.

2.1 Illumination Distribution

According to the above distant illumination assumption, we describe the intensity of the incident light ray from the direction (θ, ϕ) to a scene as $L(\theta, \phi)$. Here, θ and ϕ are the zenith and azimuth angles in the spherical coordinate system centered at the scene. Hereafter, we call $L(\theta, \phi)$ the illumination distribution of the scene.

We can represent an illumination distribution as a linear combination of basis functions. It is known that low-degree spherical harmonics can be used for the basis functions of an illumination distribution when a scene of interest is convex and obeys the Lambert model [10]. On the other hand, it is known that a large number of spherical harmonics are required when a scene is not convex and cast shadows are observed [14].

Ng *et al.* [7] show that an illumination distribution can approximately be represented by using a small number of Haar wavelets sparsely. Okabe *et al.* [9] make use of Haar wavelets for inverse lighting from cast shadows, and show that Haar wavelets work better than spherical harmonics for the basis functions of an illumination distribution. Therefore, our proposed method represents an illumination distribution by using Haar wavelets as

$$L(\theta, \phi) = \sum_{n=1}^{N} \alpha_n H_n(\theta, \phi), \tag{1}$$

where $H_n(\theta, \phi)$ and α_n are the n-th basis function of Haar wavelets and the corresponding n-th coefficient.

Based on the assumption of known shape and reflectance, we can synthesize the basis images; the n-th basis image is the image of the scene when the illumination distribution is equal to the n-th basis function of Haar wavelets $H_n(\theta, \phi)$. We denote the p-th $(p = 1, 2, 3, ..., P)$ pixel value in the n-th basis image by R_{np}. According to the superposition principle, the p-th pixel value of a single input image I_p $(p = 1, 2, 3, ..., P)$ is described as

$$I_p = \sum_{n=1}^{N} \alpha_n R_{np}. \tag{2}$$

2.2 Radiometric Response Function

The radiance values of a scene are converted to pixel values via in-camera processing. As mentioned in Section 1, the relationship between radiance values and pixel values is described by a radiometric response function. Consumer cameras

often apply nonlinear radiometric response functions for converting radiance values to pixel values in order to improve perceived image quality [3]. In general, the response function depends on cameras themselves as well as camera settings.

Let us denote a response function by f, and assume that a radiance value I is converted to a pixel value I' by using the response function as $I' = f(I)$. Since the response function is strictly increasing, there exists the inverse of f, *i.e.* an inverse response function g. The inverse response function converts a pixel value I' to a radiance value I as $I = g(I')$, and therefore has 255 degrees of freedom for 8-bit images. Such a high degree of freedom could make the joint estimation of an illumination distribution and a response function from a single image intractable.

Accordingly, we make use of the polynomial representation for a radiometric response function [5]. Specifically, our proposed method represents the inverse response function $g(I'_p)$ as

$$I_p = g(I'_p) = I'_p + I'_p(I'_p - 1)\sum_{m=1}^{M} \beta_m I_p'^{M-m}, \tag{3}$$

where $(M+1)$ is the degree of polynomial and β_m is the m-th coefficient. Both the radiance values and pixel values are normalized so that $0 \leq I_p \leq 1$ and $0 \leq I'_p \leq 1$. The polynomial representation in eq.(3) automatically satisfies the boundary conditions: $g(0) = 0$ and $g(1) = 1$. Since the inverse response function is also strictly increasing, the coefficients β_m have to satisfy

$$g\left(\frac{I'}{255}\right) < g\left(\frac{I'+1}{255}\right), \tag{4}$$

where $I' = 0, 1, 2, ..., 254$ for 8-bit images.

2.3 Joint Estimation

Substituting eq.(3) into the left-hand side of eq.(2), we obtain

$$I'_p + I'_p(I'_p - 1)\sum_{m=1}^{M} \beta_m I_p'^{M-m} = \sum_{n=1}^{N} \alpha_n R_{np}, \tag{5}$$

i.e. a single constraint on the coefficients of the illumination distribution α_n and those of the response function β_m per pixel. Therefore, our proposed method estimates the coefficients of the illumination distribution and those of the response function by minimizing the following cost function

$$\frac{1}{P}\sum_{p=1}^{P}\left[I'_p + I'_p(I'_p - 1)\sum_{m=1}^{M}\beta_m I_p'^{M-m} - \sum_{n=1}^{N}\alpha_n R_{np}\right]^2 + \frac{w}{N}\sum_{n=1}^{N}|\alpha_n|_1 \tag{6}$$

with respect to α_n and β_m subject to the strictly increasing constraint in eq.(4) and the non-negativity constraint on the illumination distribution [1]. Here, the

[1] The non-negativity constraint is represented by the linear constraints on the coefficients of the illumination distribution α_n.

first term measures the discrepancy between the radiance values predicted by the inverse response function and the illumination distribution, and the second term measures the sparseness of the coefficients for the illumination distribution. Once the coefficients of the illumination distribution α_n and those of the response function β_m are obtained, we can compute the illumination distribution and the response function by substituting them into eq.(1) and eq.(3).

2.4 Joint Estimation with Approximate Geometry

The performance of inverse lighting depends on the accuracy of the given geometry of a scene. Unfortunately, however, inverse lighting is often used with approximate geometry; it is used for a previously captured image in a passive manner where the geometry is given manually [12], and for a building block of multi-view inverse rendering where detailed shape is unknown and to be estimated [15, 13, 16, 4].

Accordingly, in order to make the joint estimation with approximate geometry robust against inaccurate geometry, we propose to mask the pixels near the edges in an input image out of the estimation. The example of the edges and mask are shown in Fig. 5. The reason why the pixels near the edges are removed is that the edges correspond to the discontinuity in depths and surface normals and the sharp shadow boundaries caused by the occluding objects, and therefore their pixel values computed from the given geometry significantly change due to slight noises in the geometry.

3 EXPERIMENTS

To confirm the effectiveness of our proposed method, we conducted qualitative and quantitative evaluation by using both synthetic and real images.

3.1 Comparison Using Synthetic Images

We compared the following four methods.

- **HW + RF (Our Proposed Method)** jointly estimates the illumination distribution of a scene and the inverse response function of a camera from a single input image. Haar wavelets are used for the basis functions of the illumination distribution.
- **HW** assumes that the inverse response function is linear and estimates only the illumination distribution by using Haar wavelets in a similar manner to the existing technique [9].
- **L-SH + RF [8]** jointly estimates the illumination distribution and the inverse response function. It assumes a convex scene, and therefore low-degree spherical harmonics up to the second degree are used for the basis functions of the illumination distribution.
- **H-SH + RF** jointly estimates the illumination distribution and the inverse response function. High-degree spherical harmonics are used.

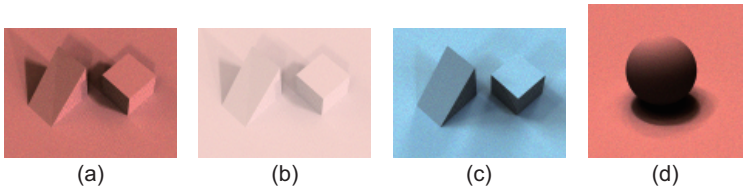

Fig. 1. The images of four scenes synthesized by using the HDR radiance maps. Each one is used as input for inverse lighting.

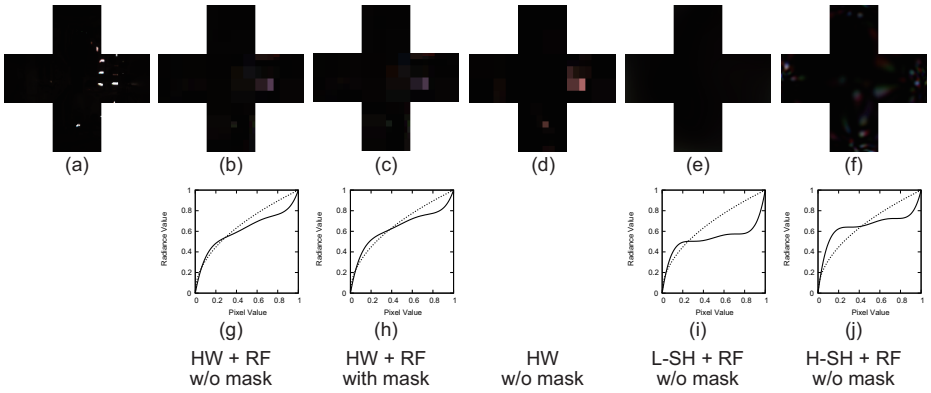

Fig. 2. The results of the first scene: (a) the ground truth of the illumination distribution, and those estimated by using (b) HW + RF w/o mask, (c) HW + RF with mask, (d) HW, (e) L-SH + RF, and (f) H-SH + RF respectively, and the corresponding inverse response functions estimated by using (g) HW + RF w/o mask, (h) HW + RF with mask, (i) L-SH + RF, and (j) H-SH + RF respectively, where the ground truth (dotted line) is superimposed.

For HW + RF and HW, the illumination distribution is represented by the Vertical Cross Cube Format (cube map) whose resolution is $N = 1280$ ($= 256 \times 5$). For L-SH + RF and H-SH + RF, the degrees of spherical harmonics are 2 and 35, *i.e.* $N = 9$ and $N = 1296$ respectively. We empirically set the degree of the polynomial to $(M + 1) = 5$. For HW + RF and HW, we used CVX [2] for optimization, and set the weight, which controls the balance between the data term and the sparseness term in eq.(6), to $w = 10^{-2}$. We show how the performance of our method depends on the value of the weight w later in Subsection 3.3.

As shown in Fig. 1, we tested four images synthesized by using the HDR radiance maps [1]. In Fig. 1 (a), two polyhedrons on a plane [2] are illuminated by the radiance map of "Galileo's Tomb", and its radiance values are converted by a convex upward inverse response function. In Fig. 1 (b), (c), and (d), one

[2] The geometry of the scene is the same as that of the real images in the next subsection.

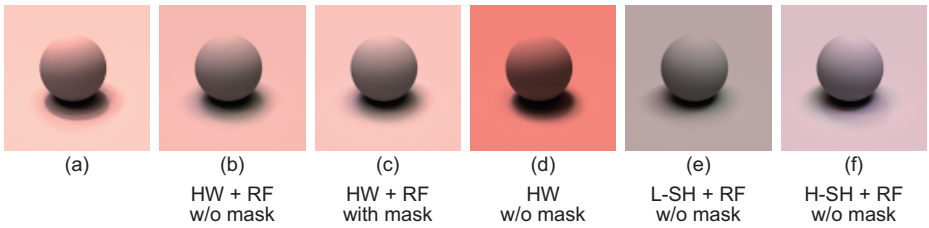

Fig. 3. The relighting images by using (a) the ground truth of the illumination distribution, and those estimated by using (b) HW + RF w/o mask, (c) HW + RF with mask, (d) HW, (e) L-SH + RF, and (f) H-SH + RF respectively.

Table 1. The RMSEs of the relighting images in Fig. 3 and those of the estimated inverse response function in Fig. 2 for the four scenes shown in Fig. 1 (a), (b), (c), and (d).

method	mask	relighting image				response function			
		(a)	(b)	(c)	(d)	(a)	(b)	(c)	(d)
HW + RF	w/o	0.056	**0.055**	0.051	**0.077**	0.077	**0.078**	**0.067**	**0.019**
HW + RF	w	**0.031**	0.101	0.063	0.079	**0.070**	0.091	0.110	0.021
HW	w/o	0.233	0.164	0.094	0.186	-	-	-	-
L-SH + RF	w/o	0.176	0.177	**0.048**	0.085	0.186	0.247	0.118	0.034
H-SH + RF	w/o	0.093	0.151	0.097	0.422	0.121	0.580	0.210	0.176

of the conditions is different from (a); (b) a convex downward inverse response function, (c) the radiance map of "Eucalyptus Grove", and (d) a sphere on a plane are used instead respectively. We added zero-mean Gaussian noises, whose standard deviation is $\sigma = 0.01$ for pixel values normalized to $[0, 1]$, to those synthetic images.

Fig. 2 shows the results of the first scene: (a) the ground truth of the illumination distribution, and those estimated by using (b) HW + RF w/o mask, (c) HW + RF with mask, (d) HW, (e) L-SH + RF, and (f) H-SH + RF respectively. We show the corresponding inverse response functions estimated by using HW + RF w/o mask, HW + RF with mask, L-SH + RF, and H-SH + RF in (g), (h), (i), and (j) respectively, where the ground truth (dotted line) is superimposed on each result. In order to evaluate the effectiveness of the estimated illumination distribution for applications to AR and MR, we synthesized the relighting images of a scene different from that of the input image. Fig. 3 shows the relighting images by using (a) the ground truth of the illumination distribution, and those estimated by using (b) HW + RF w/o mask, (c) HW + RF with mask, (d) HW, (e) L-SH + RF, and (f) H-SH + RF respectively. Table 1 summarizes the RMSEs of the relighting images in Fig. 3 and those of the estimated inverse response function in Fig. 2 for the four scenes shown in Fig. 1 (a), (b), (c), and (d).

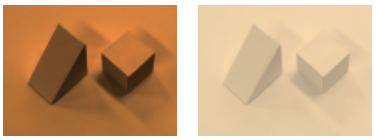

Fig. 4. The real input images of a scene captured with different response functions.

Fig. 5. The edges and mask for the real image in Fig. 4 (a).

Effectiveness of Mask:
Comparing the results of HW + RF, *i.e.* our proposed method, without and with mask in Figs. 2, 3, and Table 1, we can see that the proposed method without mask works better than that with mask. This is because the geometry of a scene is accurate for synthetic images, and therefore the use of the mask would lose the clues for estimating the high-frequency details of illumination distributions. Hence, we did not use the mask for evaluating the performance of the other methods: HW, L-SH + RF, and H-SH + RF. It is worth noting that the degradation of the performance due to the mask is relatively small.

Effectiveness of Joint Estimation:
Comparing the results of HW + RF, *i.e.* our proposed method with those of HW in Figs. 2, 3, and Table 1, we can see that our method works better than HW both qualitatively and quantitatively. In particular, the brightness of the relighting image in Fig. 3 (d) is distorted due to the assumption of a linear inverse response function. Those results show that the effectiveness of the joint estimation of an illumination distribution and an inverse response function. The reason why the estimated inverse response function in Fig. 2 (g) deviates from the ground truth when pixel values are large is that the number of bright pixels is small in the input image in Fig. 1 (a), and therefore the constraints from those pixels are loose.

Effectiveness of Haar Wavelets:
Comparing the results of HW + RF, *i.e.* our proposed method with those of L-SR + RF [8] in Figs. 2, 3, and Table 1, we can see that our method works better than L-SH + RF both qualitatively and quantitatively. Moreover, our method works better than H-SH + RF, *i.e.* even when we increase the degree of spherical harmonics from 2 (N=9) to 35 (N=1296). Specifically, the RMSEs of our method using Haar wavelets are significantly smaller than those of the

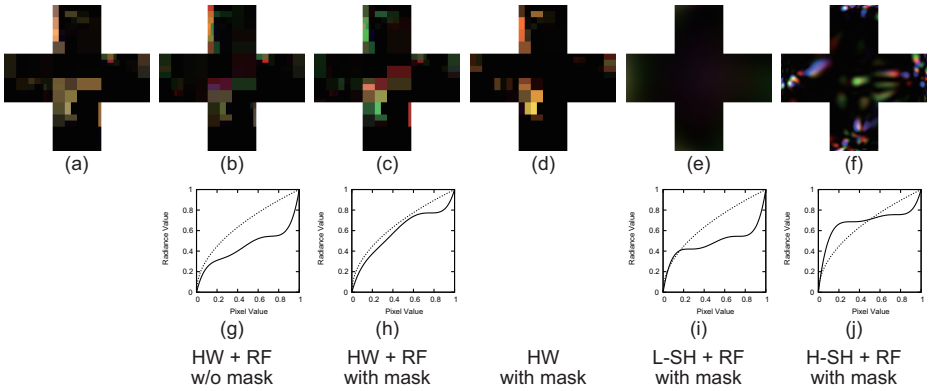

Fig. 6. The results for the first real image: (a) the ground truth of the illumination distribution, and those estimated by using (b) HW + RF w/o mask, (c) HW + RF with mask, (d) HW, (e) L-SH + RF, and (f) H-SH + RF respectively, and the corresponding inverse response functions estimated by using (g) HW + RF w/o mask, (h) HW + RF with mask, (i) L-SH + RF, and (j) H-SH + RF respectively, where the ground truth (dotted line) is superimposed.

methods using spherical harmonics. Those results show that Haar wavelets are effective for representing illumination distributions.

Hence, we can conclude that our proposed method works better than the existing/closely-related methods. In particular, we confirmed the effectiveness of both the joint estimation and the Haar wavelets.

3.2 Comparison Using Real Images

In order to show the effectiveness of our proposed method for real images, we compared our method with the other three methods by using real images in a similar manner to the previous subsection. Fig. 4 shows the real input images of a scene captured by the same Point Grey camera (Chameleon) with different settings in response functions. Fig. 5 shows (a) the edges extracted by using the Canny edge detector and (b) the mask for the real input image in Fig. 4 (a).

Fig. 6 shows the results for the first real image: (a) the ground truth of the illumination distribution, and those estimated by using (b) HW + RF w/o mask, (c) HW + RF with mask, (d) HW, (e) L-SH + RF, and (f) H-SH + RF respectively. Here, we consider the illumination distribution estimated from the image with a linear response function by using HW with mask as the ground truth. We show the corresponding inverse response functions estimated by using HW + RF w/o mask, HW + RF with mask, L-SH + RF, and H-SH + RF in (g), (h), (i), and (j) respectively, where the ground truth (dotted line) is superimposed on each result. Fig. 7 shows the relighting images by using (a) the ground truth of the illumination distribution, and those estimated by using (b) HW + RF w/o mask, (c) HW + RF with mask, (d) HW, (e) L-SH + RF, and

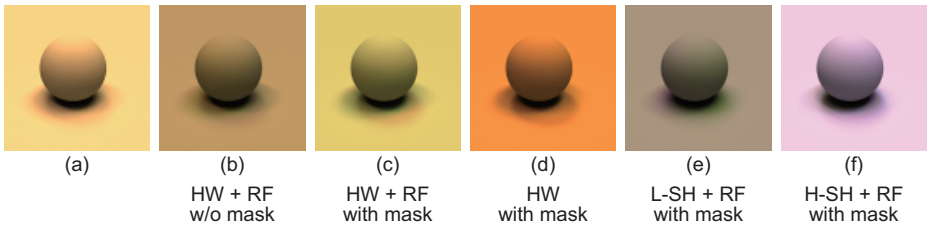

Fig. 7. The relighting images by using (a) the ground truth of the illumination distribution, and those estimated by using (b) HW + RF w/o mask, (c) HW + RF with mask, (d) HW, (e) L-SH + RF, and (f) H-SH + RF respectively.

Table 2. The RMSEs of the relighting images in Fig. 7 and those of the estimated inverse response function in Fig. 6 for the two images shown in Fig. 4 (a) and (b).

method	mask	relighting image		response function	
		(a)	(b)	(a)	(b)
HW + RF	w/o	0.206	0.036	0.231	0.095
HW + RF	w	**0.082**	**0.031**	**0.080**	**0.033**
HW	w	0.229	0.219	-	-
L-SH + RF	w	0.235	0.201	0.217	0.256
H-SH + RF	w	0.247	0.280	0.125	0.549

(f) H-SH + RF respectively. Table 2 summarizes the RMSEs of the relighting images in Fig. 7 and those of the estimated inverse response function in Fig. 6 for the two real images shown in Fig. 4 (a) and (b).

Effectiveness of Mask:

Comparing the results of HW + RF, *i.e.* our proposed method, without and with mask in Figs. 6, 7, and Table 2, we can see that the proposed method with mask works significantly better than that without mask. Those results show the effectiveness of the mask when estimating from real images with approximate geometry. Therefore, we used the mask for evaluating the performance of the other methods: HW, L-SH + RF, and H-SH + RF.

Effectiveness of Joint Estimation:

Comparing the results of HW + RF, *i.e.* our proposed method with those of HW in Figs. 6, 7, and Table 2, we can see that our method works better than HW both qualitatively and quantitatively. Those results are consistent with the results for synthetic images. The reason why the estimated inverse response function in Fig. 6 (h) deviates from the ground truth when pixel values are large is that the number of bright pixels is small in the input image in Fig. 4 (a), and therefore the constraints from those pixels are loose.

Effectiveness of Haar Wavelets:

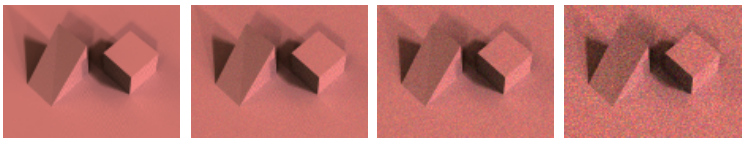

Fig. 8. The input images for the sensitivity analysis. We added zero-mean Gaussian noises with the standard deviation (a) $\sigma = 0$, (b) $\sigma = 0.01$, (c) $\sigma = 0.02$, and (d) $\sigma = 0.04$ respectively.

Table 3. The RMSEs of the relighting images synthesized from the illumination distributions estimated for various combinations of σ and w.

$\sigma \backslash w$	0	10^{-4}	10^{-3}	10^{-2}	10^{-1}
0	**0.012**	**0.012**	0.018	0.106	0.399
0.01	0.134	0.123	0.071	**0.056**	0.383
0.02	0.175	0.170	0.151	**0.065**	0.260
0.04	0.203	0.197	0.178	0.155	**0.064**

Comparing the results of HW + RF, *i.e.* our proposed method with those of L-SR + RF [8] in Figs. 6, 7, and Table 2, we can see that our method works better than L-SH + RF both qualitatively and quantitatively. Moreover, our method works better than H-SH + RF, *i.e.* even when we increase the degree of spherical harmonics from 2 (N=9) to 35 (N=1296). Those results are also consistent with the results on synthetic images.

Hence, we can conclude that our proposed method works better than the existing/closely-related methods. In particular, we confirmed the effectiveness of the mask for real images with approximate geometry in addition to the effectiveness of the joint estimation and the Haar wavelets.

3.3 Sensitivity Analysis

The performance of our proposed method depends on the value of the weight w in eq.(6) as well as noises. Accordingly, we studied the sensitivity of our method to the value of the weight w and noises by using synthetic images. Specifically, we added zero-mean Gaussian noises with varying standard deviation σ to the synthesized image of the first scene as shown in Fig. 8, and then estimated the illumination distribution and the inverse response function with varying weight w. Table 3 summarizes the RMSEs of the relighting images synthesized from the illumination distributions estimated for various combinations of σ and w.

We can see that the sparseness term, *i.e.* non-zero weight w is not always necessary when using noiseless image in Fig. 8 (a). On the other hand, we can see that the sparseness term works well when Gaussian noises are added to the input image. In particular, the performance of our proposed method is not sensitive to

the value of the weight, and $w = 10^{-2}$ is the best for realistic noise levels such as $\sigma = 0.01$ and $\sigma = 0.02$.

4 CONCLUSION

In this paper, we proposed a method for jointly estimating both the illumination distribution of a scene and the radiometric response function of a camera from cast shadows in a single input image. Specifically, our proposed method represents the illumination distribution and the response function by using Haar wavelets with the sparseness constraint and polynomials respectively, and then estimates their coefficients. We conducted a number of experiments using both synthetic and real images, and confirmed that our method works better than the existing methods. Incorporating our method into multi-view inverse rendering [15, 13, 16, 4] is one of the directions of our future work.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP17H00744.

References

1. Debevec, P.: Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In: Proc. ACM SIGGRAPH1998. pp. 189–198 (1998)
2. Grant, M., Boyd, S.: CVX: Matlab software for disciplined convex programming, version 2.1. http://cvxr.com/cvx (2014)
3. Grossberg, M., Nayar, S.: What is space of camera response functions? In: Proc. IEEE CVPR2003. pp. 602–609 (2003)
4. Kim, K., Torii, A., Okutomi, M.: Multi-view inverse rendering under arbitrary illumination and albedo. In: Proc. ECCV2016 (LNCS 9907). pp. 750–767 (2016)
5. Lee, J.Y., Matsushita, Y., Shi, B., Kweon, I.S., Ikeuchi, K.: Radiometric calibration by rank minimization. IEEE TPAMI **35**(1), 144–156 (2013)
6. Marschner, S., Greenberg, D.: Inverse lighting for photography. In: Proc. IS&T/SID Fifth Color Imaging Conference. pp. 262–265 (1997)
7. Ng, R., Ramamoorthi, R., Hanrahan, P.: All-frequency shadows using non-linear wavelet lighting approximation. In: Proc. ACM SIGGRAPH2003. pp. 376–381 (2003)
8. Ohta, S., Okabe, T.: Does inverse lighting work well under unknown response function? In: Proc. VISAPP2015. pp. 652–657 (2015)
9. Okabe, T., Sato, I., Sato, Y.: Spherical harmonics vs. Haar wavelets: basis for recovering illumination from cast shadows. In: Proc. IEEE CVPR2004. pp. I–50–57 (2004)
10. Ramamoorthi, R., Hanrahan, P.: A signal-processing framework for inverse rendering. In: Proc. ACM SIGGRAPH2001. pp. 117–128 (2001)
11. Sato, I., Sato, Y., Ikeuchi, K.: Acquiring a radiance distribution to superimpose virtual objects onto a real scene. IEEE TVCG **5**(1), 1–12 (1999)

12. Sato, I., Sato, Y., Ikeuchi, K.: Illumination from shadows. IEEE TPAMI **25**(3), 290–300 (2003)
13. Shan, Q., Adams, R., Curless, B., Furukawa, Y., Seitz, S.: The visual Turing test for scene reconstruction. In: Proc. IEEE 3DV. pp. 25–32 (2013)
14. Sloan, P.P., Kautz, J., Snyder, J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In: Proc. ACM SIGGRAPH2002. pp. 527–536 (2002)
15. Wu, C., Wilburn, B., Matsushita, Y., Theobalt, C.: High-quality shape from multi-view stereo and shading under general illumination. In: Proc. IEEE CVPR2011. pp. 969–976 (2011)
16. Zollhöfer, M., Dai, A., Innman, M., Wu, C., Stamminger, M., Theobalt, C., Nießner, M.: Shading-based refinement on volumetric signed distance functions. ACM TOG **34**(4), 1–14 (2015)