Determining the Radiometric Response Function from a Single Grayscale Image

Stephen Lin
Microsoft Research Asia

Lei Zhang
Peking University*

Abstract

A method is presented for computing the radiometric response function of a camera from a single grayscale image. While most previous techniques require a set of registered images with different exposures to obtain response data, our approach capitalizes on a statistical feature of graylevel histograms at edge regions to gain information for radiometric calibration. Appropriate edge regions are automatically determined by our technique, and a prior model of radiometric response functions is employed to deal with incomplete data. With this single-image method, radiometric calibration becomes possible to perform in many instances where the camera is unknown.

1 Introduction

In most cameras, there exists a nonlinear relationship between the image intensities output from the camera and the scene radiance that enters the imaging system. This nonlinearity is intentionally designed into digital cameras to compress the dynamic range of scenes and to account for nonlinearities in display systems. While this feature may be beneficial for viewing purposes, it impairs many computer vision methods which assume that image intensities are linearly related to scene radiance.

A broad range of vision algorithms require this linearity because they need precise measurements of scene radiance for accurate processing. In photometric methods such as shape from shading, color constancy and illumination estimation, physical information is derived from scene radiance for analyzing a scene. Image intensities are also implicitly presumed to convey scene radiance in many other vision methods such as object recognition and multi-camera stereo, so that images captured from different cameras or at different brightness levels can be properly compared. To obtain scene radiance information from images, the mapping from scene radiance to image intensity, called the radiometric response function, must be determined to undo its nonlinear effects on image intensities.

*This work was done while Lei Zhang was an intern at Microsoft Research Asia.

1.1 Background

In the imaging process, radiance from a scene first is recorded at the imaging array of a camera as image irradiance. These irradiance values I are then transformed according to the camera's radiometric response f into measured intensities M that are output from the camera:

$$M = f(I).$$

For vision algorithms that require irradiance values as input, an inverse response function $g = f^{-1}$ should be recovered to map image intensities into irradiance values, in a process known as radiometric calibration. Since sensor output increases monotonically with respect to I, response functions f are invertible.

A mapping g can be determined by fitting a function to corresponding values of intensity and irradiance. For intensity values in measured images, however, their corresponding irradiance values are generally unknown. Radiometric calibration methods traditionally deal with this problem by varying the camera exposure in a set of images, which effectively modulates the irradiance captured at the imaging array. For a pair of images, the ratio of captured radiance at corresponding pixels is equal to the exposure ratio between the two images, as expressed in the following relationship:

$$g(m_A) = kg(m_B) \tag{1}$$

where m_A denotes measured image intensities in image A, m_B represents intensities of corresponding points in image B, and k denotes the exposure ratio between A and B. From a set of corresponding points in the image pair, an inverse response function g can be solved using Eq. (1), with the assumption that the sensor response does not change over the image grid, such as from vignetting or fixed pattern noise.

1.2 Previous work

Most previous works require that the image sequence with varying exposure be captured by a fixed camera. For

instances where the ratios of the exposure values are known, Mann and Picard [9] compute a parametric response function in the form of a gamma curve, and Debevec and Malik [1] employ a smoothness constraint to obtain a non-parametric response function. The need for precise exposure ratios is avoided in the method by Mitsunaga and Nayar [10], which iteratively solves for a polynomial response function beginning with just a rough estimate of exposure ratios. Iterative methods have also been proposed by Tsin et al. [14], which estimates non-parametric responses using a statistical model of the CCD imaging process, and by Pal et al. [12], which utilizes probabilistic imaging models and prior models of response functions to compute radiometric response functions that can differ among the images in the sequence.

Less restrictive on the input image sequence, a few previous methods permit some camera movement or scene motion, with changes in exposure level. Mann [8] presents an iterative method for response function estimation from sequences captured by a rotating and zooming camera. Kim and Pollefeys [6] propose a technique that computes point correspondences to allow free movement of the camera and some motion in the scene. The need for spatial correspondences is circumvented by Grossberg and Nayar [4], who relate histogrammed intensity values between two images of different exposure.

In many scenarios such as web image analysis, the image capture device is not accessible, so multiple images of different exposure cannot be obtained for radiometric calibration. In such cases, single-image approaches are needed to estimate the response function. Nayar and Mitsunaga [11] introduce a method for increasing the dynamic range of digital cameras by placing upon the imaging array an optical filter with spatially varying transmittance. For images captured with this filter, the response function can be computed from neighboring pixels of different exposure and constant scene brightness. Farid [2] presents a technique for estimating the gamma correction of an image from certain frequency-domain correlations that it introduces. However, this approach makes assumptions on the statistics of scene radiance, and the radiometric responses of many cameras differ significantly from a gamma function.

A more general single-image method is proposed by Lin et al. [7], which obtains information about the radiometric response from color distributions of local edge regions. Due to blending of distinct region colors, irradiance colors from edge regions should form linear distributions in color space. But because of nonlinear radiometric response functions, measured edge colors actually compose nonlinear distributions that are directly related to the response function. With a prior model of response functions compiled by Grossberg and Nayar [5], the inverse radiometric response is computed as the function that maps the nonlinear distributions of measured edge colors into linear distributions.

This color-based approach, however, cannot be applied to grayscale images, which are popular in artistic photography and are commonly used in vision applications such as face recognition and tracking in surveillance video. Unlike color which is commonly measured as a vector of three spectral elements, grayscale intensities are scalar quantities for which properties such as linearity and nonlinearity do not apply. To address the problem of radiometric calibration from a single grayscale image, we propose in this paper a 1D analogue of the 3D color method in [7].

In this work, we similarly consider the formation of edge irradiance distributions from region blending. But since a nonlinear response function does not transfigure the locus of 1D grayscale values as it does for 3D colors, we propose to use higher-order distribution features along image edges to extract radiometric response information. For an edge pixel that views two scene regions as shown in Fig. 1(a), the ratio of the two region areas that it observes can be considered a uniform random variable, and consequently the grayscale histogram of edge irradiance values should statistically be uniform between the graylevels of the two bounding regions, as displayed in Fig. 1(b). In contrast, the edge pixel intensities in a measured image form a skewed histogram distribution, as illustrated in Fig. 1(d), that is directly related to the nonlinearity of the radiometric response function, displayed in Fig. 1(c). From this property, we show that the radiometric response can be computed as the function that transforms the non-uniform histograms of edge regions into uniformly distributed histograms. Our experimentation has shown this technique to obtain accurate response functions from single grayscale images.

2 Histogram Analysis

Our proposed method derives the radiometric response function from histograms of measured intensities at edge regions. In this section, we describe how region blending at edges leads to statistical uniformity of irradiance histograms, and how sensor nonlinearities map these irradiance distributions into non-uniform histograms of measured intensities output by the camera. From these measured intensity histograms, the inverse response function can be determined as the transformation that unskews their distributions.

2.1 Histograms of Image Irradiance

To explain the uniformity of irradiance histograms, we first review the image sensing process. Images are formed on a CCD sensor array that records radiance from the scene. Because the array is limited in spatial resolution, each array element $\boldsymbol{x}$ images a solid angle of the scene, and we denote

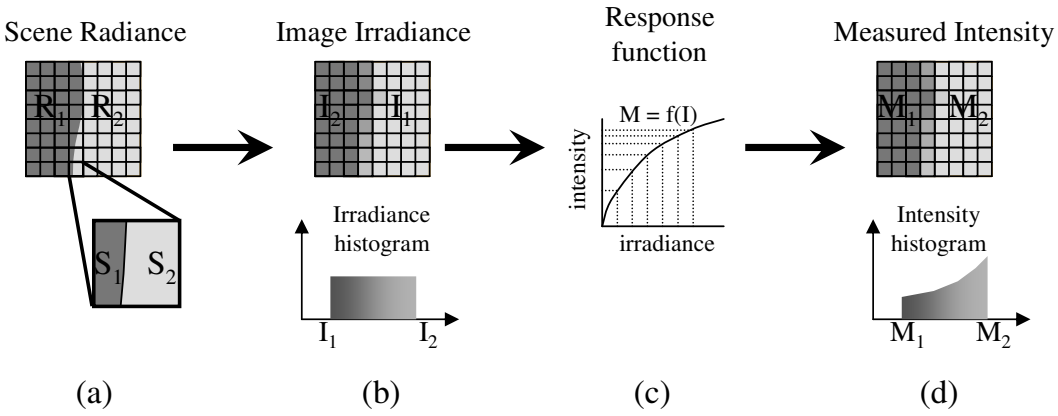

Figure 1. Transformation of scene radiance to measured intensities in local edge regions. (a) An edge bounded by regions with scene radiance R_1 and R_2. A pixel that straddles the edge receives radiance from both regions in proportion to areas S_1 and S_2 that it images of each region. (b) Each imaging element records a single grayscale value as image irradiance. Since the proportion between S_1 and S_2 in an edge pixel is a uniform random variable, a histogram of irradiance values is uniform (not including irradiance values of bounding regions). (c) A camera response f nonlinearly maps irradiance values to measured intensities. (d) Because of the nonlinear mapping, the histogram of measured intensity values has a non-uniform distribution.

the set of image plane points within an array element as $S(\boldsymbol{x})$.

The image irradiance I at $\boldsymbol{x}$ depends on the sensitivity q of the array element with respect to wavelength λ and the incoming scene radiances R incident upon image plane points p in $S(\boldsymbol{x})$:

$$I(\boldsymbol{x}) = \int_{\lambda} \int_{p \in S(x)} R(p, \lambda) q(\lambda) \; dp \; d\lambda. \tag{2}$$

For an edge that separates two scene regions of distinct but uniform colors, an edge pixel element generally views portions of both bounding regions, as illustrated in Fig. 1(a). For an edge pixel $\boldsymbol{x}$ divided into regions $S_1(\boldsymbol{x})$ and $S_2(\boldsymbol{x})$ with respective scene radiances $R_1(\lambda)$ and $R_2(\lambda)$, the overall radiance incident on $\boldsymbol{x}$ can be expressed as

$$\int_{p \in S(x)} R(p, \lambda) dp = \int_{p \in S_1(x)} R_1(\lambda) dp + \int_{p \in S_2(x)} R_2(\lambda) dp$$

$$= \alpha R_1(\lambda) + (1 - \alpha) R_2(\lambda) \tag{3}$$

where $\alpha = \int_{p \in S_1(x)} dp$ and $S(\boldsymbol{x})$ is of unit area.

Substituting Eq. (3) into Eq. (2), we can relate the image irradiance of $\boldsymbol{x}$ to the irradiance values of the bounding regions:

$$I(\boldsymbol{x}) = \alpha \int_{\lambda} R_1(\lambda) q(\lambda) d\lambda + (1 - \alpha) \int_{\lambda} R_2(\lambda) q(\lambda) d\lambda$$

$$= \alpha I_1(\boldsymbol{x}) + (1 - \alpha) I_2(\boldsymbol{x})$$

An edge window typically contains numerous edge pixels composed of varying proportions (α, $1 - \alpha$) of the two bounding regions. Since these proportions are determined by chance according to edge paths, α can be considered a uniform random variable in the range $(0, 1)$. With this property, a histogram of image irradiance values should be statistically uniform in an interval between the irradiances of the bounding regions, as shown in Fig. 1(b). Since a window will also contain many non-edge pixels with α equal to zero or one, this histogram uniformity is valid in practice over a contracted interval that excludes the graylevels of the bounding regions and their image noise.

2.2 Histograms of Measured Intensities

From the imaging array to the output image, the irradiance values undergo transformations represented by the camera's radiometric response function. This mapping from irradiance to measured intensity is illustrated in Fig. 1(c). Because of the nonlinearity of the response function, the irradiances are converted into intensity values that form a non-uniform histogram, as exemplified in Fig. 1(d). Our method takes advantage of these histogram non-uniformities to determine response functions.

This edge characteristic exploited for radiometric calibration differs from that used in the color-based method of [7]. The color method takes advantage of how region blending constrains the *values of irradiance* to form linear distributions in RGB space and how a response function transforms these values into a nonlinear configuration. This sort of nonlinear shape deformation cannot occur for scalar grayscale values, which are confined to a 1D line. To obtain information on the response function, our technique capitalizes on a statistical feature of the *frequency of irradiance values* that arises from region blending.

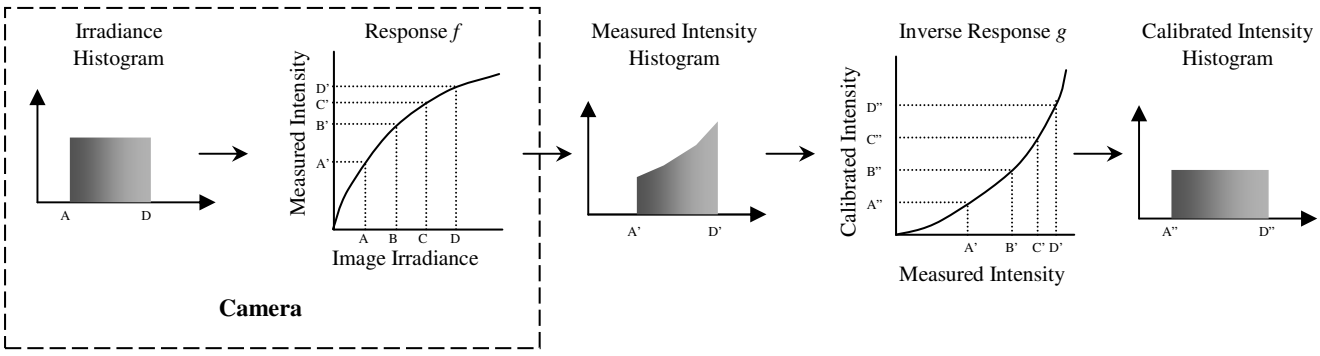

Figure 2. The inverse response function g should transform a non-uniform histogram of measured intensities at an edge region into a uniform histogram of calibrated intensities.

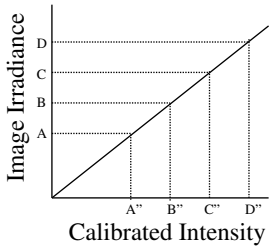

Figure 3. Composition of inverse response function g and response function f in Fig. 2 gives a linear relationship between calibrated intensity values and image irradiance.

2.3 Histogram Transformations

To determine the inverse response function from measured intensities in edge regions, our method computes the function that maps non-uniform histograms of measured intensities into uniform distributions of calibrated intensities, as illustrated in Fig. 2. As shown in Fig. 3, the composition of the inverse response function g with the response function f yields calibrated intensity values $g(f(I))$ that are linearly related to irradiance. To deal with the scale difference between irradiance and calibrated intensity, we normalize both the domain and co-domain of the function g so that it satisfies $g(0) = 0$ and $g(1) = 1$, as done in [5].

To evaluate the uniformity of a histogram H, we utilize a distance function that is an approximate measure of histogram non-uniformity:

$$N(H) = \frac{1}{b} \sum_{k=I_{min}}^{I_{max}} \left[|H(k)| - \frac{|H|}{b} \right]^2 + \frac{\beta}{3} \sum_{n=1}^{3} \left[|H_n| - \frac{|H|}{3} \right]^2$$

$$\text{where } |H_n| = \sum_{i=I_{min}+\frac{(n-1)b}{3}}^{I_{min}+\frac{nb}{3}-1} |H(i)| \qquad (4)$$

where $|H|$ represents the number of pixels recorded in H, $|H(k)|$ denotes the number of pixels of intensity k, $|H_n|$ stands for the number of pixels in intensity range n, I_{min} is the minimum intensity in H, I_{max} is the maximum intensity in H, $b = I_{max} - I_{min} + 1$ denotes the number of graylevels in H, and β is an empirically fixed coefficient. I_{min} and I_{max} are set to exclude the graylevels and noise of the bounding regions, and their range is restricted to be a multiple of 3, so that the three histogram ranges H_n are equal. Transforming a histogram by an inverse response function leads to real-valued graylevels $p.q$, where p denotes the integer component and q the fractional component. These graylevels are distributed between two histogram bins such that $|H(p)|$ is incremented by $(1 - q)$ and $|H(p + 1)|$ by q.

This measure combines the variance of two components: the number of elements in each graylevel, and the number of elements in each of three grayscale bins that each cover a third of the histogram range. While the first term quantifies unevenness among individual graylevels, the second term is a measure of unbalance among different sections of the distribution. We note that entropy measures do not specifically evaluate distribution balance, which is important for our algorithm. The Earth Movers Distance [13] computed with respect to an even distribution of a given number of samples in a given number of bins could alternatively be used as a more precise metric of histogram non-uniformity, but at a greater computational cost.

With this measure, the inverse response function g is computed from the histograms of all the edge regions obtained from an image, as illustrated in Fig. 4. The benefits of using a larger number of edge regions are to cover a broader range of image intensity values and to reduce the influence of noise. For a given image, the intensity histograms of the edge regions are gathered into an observation set Ω, and the

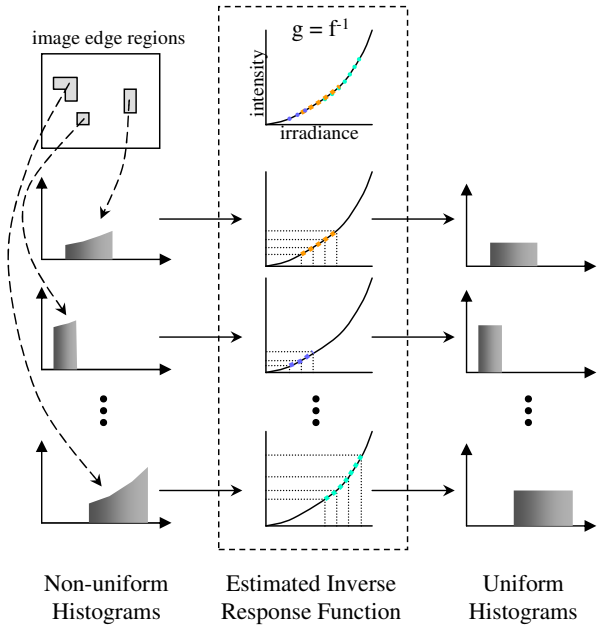

Figure 4. Estimation of inverse response function g as the function that transforms the non-uniform histograms of edge regions into uniform histograms.

total distance of a solution g is given as

$$D(g;\Omega) = \sum_{H \in \Omega} w_H N(g(H)), \tag{5}$$

where $w_H = |H|/b$ is a coefficient that gives greater weight to more densely populated histograms that should have more statistically accurate distributions. The desired inverse response function g should yield the smallest total distance.

3 Selection of Edge Regions

The edge histograms in Ω are tabulated from image regions that exhibit intensity blending between two areas of uniform graylevels. To form these edge regions, our method first identifies non-overlapping, fixed-size (15×15) image windows that have the blending characteristics required by our approach. Valid windows are centered on a Canny-detected edge whose path divides the window into exactly two regions. The edge path is dilated by three pixels, and the intensity variances of the two partitioned non-edge regions are computed. If the variances of both regions lie below a specified threshold, then the regions are considered uniform. Additionally, the edge intensities must lie between the graylevels of the two bounding regions, because of the monotonicity of response functions. This requirement helps to exclude some edges that exhibit artifacts such as JPEG ringing and chromatic aberration.

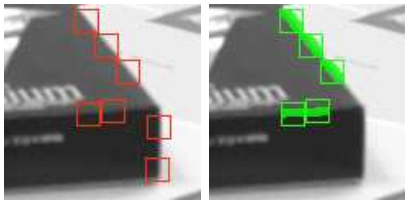

Figure 5. Formation of edge regions, shown in a zoomed image area. (a) Valid windows; (b) Aggregation of valid windows into two edge regions. Two windows containing vertical edges were excluded because of anomalous intensity histograms. The compression artifacts are due to pdf.

For greater robustness of the statistical non-uniformity measure, valid edge windows are aggregated into larger edge regions to increase the number of pixel samples in a histogram. This is performed by fusing adjacent or approximately adjacent windows that have corresponding region intensities, as exemplified in Fig. 5.

There exist some cases, such as perfectly vertical or horizontal edges with respect to the CCD array, where an edge region may have an irradiance distribution that is not uniform. Undesirable edge regions caused by chromatic aberration may also be inadvertantly included in the observation set, since the presence of this artifact is sometimes not apparent in a grayscale image. To reduce the number of such unfavorable histograms in the observation set, our method rejects edge regions whose intensity histograms do not exhibit a skew that is characteristic of most response functions. Since practically all common response functions skew uniform irradiance histograms towards brighter intensities, we detect outlier intensity histograms as those that do not generally display this form, which is roughly expressed as $|H_1| \leq |H_2| \leq |H_3|$ where $|H_n|$ is defined in Eq. (4).

4 Response Function Estimation

The accuracy of an inverse response function estimated from the distance function of Eq. (5) is limited by the amount of information present in Ω, whose histograms may not fully span the range of possible intensity values. To deal with missing data, physical characteristics of response functions such as smoothness and monotonicity could potentially be used to constrain the solution. In this work, we instead employ more focused information from the DoRF database of real-world response functions compiled by Grossberg and Nayar [5].

The principal components of the DoRF database are used to represent inverse response functions g, as described in [5]:

$$g = g_0 + H\,c \tag{6}$$

where g_0 is the mean inverse response, and H is a matrix

whose columns are composed of the first $N = 5$ eigenvectors. c is an N-dimensional coefficient vector that represents an inverse response function g.

Our algorithm further employs this prior data to regulate the distance function of Eq. (5) in a MAP estimation framework. A prior model $p(g)$ of inverse response functions is formed from the DoRF database as a Gaussian mixture model:

$$p(g) = \sum_{i=1}^{K} \alpha_i \mathcal{N}(g; \mu_i, \Sigma_i). \quad (7)$$

In our implementation, we empirically use nine kernels ($K = 9$) obtained using the EM algorithm.

The inverse response function g should yield a low total distance as expressed in Eq. (5), so we model a likelihood function $p(\Omega|g)$ by incorporating this distance measure into an exponential distribution:

$$p(\Omega|g) = \frac{1}{Z} exp(-\lambda D(g; \Omega))$$

where λ is set empirically to 10^4 and Z is a normalization constant.

For the observation set Ω of a given image, the MAP estimate of the response function g^* can then be expressed in terms of the prior $p(g)$ and the likelihood $p(\Omega|g)$:

$$g^* = \arg\max p(g|\Omega) = \arg\max p(\Omega|g)p(g).$$

By taking the $\log$ of $p(g|\Omega)$, g^* also can be written as

$$g^* = \arg\min E(g) = \arg\min \lambda D(g; \Omega) - \log p(g),$$

where g^* is the optimal solution of the objective function $E(g)$.

The optimization is computed by the Levenberg-Marquardt method using nine initial values of g that correspond to the Gaussian centers μ_i of the GMM. Since g is represented by principal components in Eq. (6), the first and second derivatives of $g(c)$ are approximated by the first and second differences with a small δc. After the optimization algorithm converges, the result is refined sequentially in each dimension using a greedy local search.

5 Results

To evaluate our technique, it was applied on sets of single-image input from a grayscale video camera (Ikegata IK-110) and a color SLR camera (Canon EOS-1Ds) whose RGB readings were converted to grayscale by $I = 0.30R + 0.59G + 0.11B$ [3]. In the case of grayscale values converted from color images, it is assumed that RGB demosaicking does not introduce any bias to edge intensity distributions. None of these cameras are included in the DoRF database used to form the prior model.

Camera	RMSE	Disparity
CANON EOS-1Ds	0.0117	0.0216
Ikegata IK-110	0.0194	0.0490

Table 1. RMSE and Disparity of Estimated Inverse Response Functions

For each camera, experiments were performed on 20 individual images of various scenes, and representative results are displayed in Fig. 6 and Fig. 7. For the edge regions of an image, the ranges of its histograms can be plotted to give an indication of the completeness of the data. The response charts include the estimated inverse response function of the given image, the most inaccurate response function estimated from among the 20-image set, the inverse response function recovered using all 20 images together, and measurements from a Macbeth ColorChecker chart. For the Canon EOS-1Ds, an additional response function is computed using the multiple-exposure method in [10]. This comparison is not included for the Ikegata camera, since its exposure level cannot be adjusted.

A quantitative comparison of our method to the other techniques is tabulated in Table 1 over the image sets. The RMS error is computed in terms of normalized irradiance, and the disparity represents the average value among the images of the maximum difference between our function and the comparison curve. For the Canon EOS-1Ds, the comparison is made to the average response function using the method in [10] and the Macbeth ColorChecker, whose measurements are interpolated using a fifth degree polynomial. For the grayscale camera, the ColorChecker alone is used for comparison. The low RMSE and disparity values demonstrate an agreement between our calibration method and others that utilize greater information. The difference in performance between the images of the two cameras can primarily be attributed to the lower resolution of the Ikegata IK-110, which resulted in less calibration data.

6 Discussion

With more complete amounts of edge data, there is less reliance on the prior model to determine a most likely response function for the given data. In contrast to the color-based method of [7] in which the completeness of data is related to its coverage of the 3D RGB space, our grayscale method needs data that extends over a 1D grayscale range. Greater quantities of data in the range, however, are needed for more statistical robustness.

For a single grayscale image, a potential method for obtaining a broader range and greater quantity of histogram data is to more aggressively acquire edge regions, such as by allowing variation in window sizes and less restrictive selection thresholds. This is a direction we plan to investi-

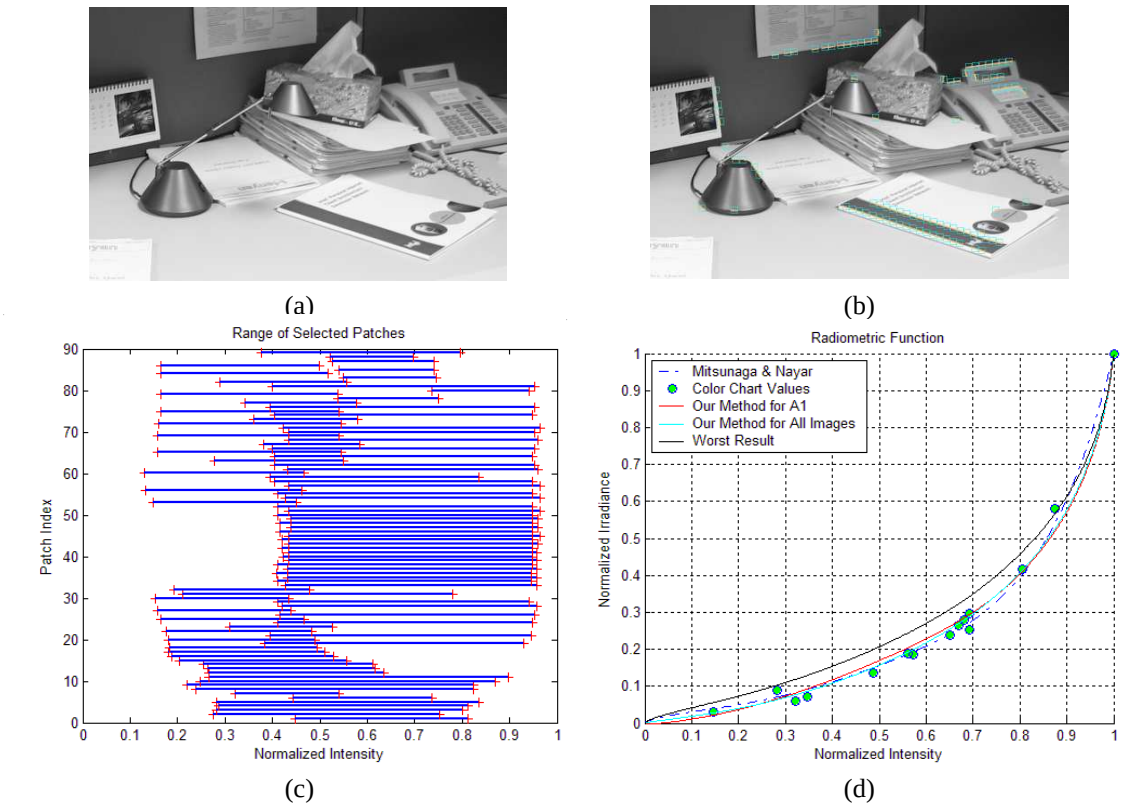

Figure 6. Results for a representative test image from a Canon EOS-1Ds camera with output converted to grayscale. (a) Test image; (b) Detected edge patches; (c) Grayscale ranges of edge regions; (d) Inverse response functions.

gate in future work. The amount of histogram data can also be increased by collecting edge regions from additional images captured by the same camera, even if the images are not registered or otherwise corresponded.

When the histograms of an image collectively span only a narrow range of intensities, the estimated inverse response function may have significant inaccuracies for graylevels outside this range. Within this range, calibrated intensities should nevertheless be linearly related to irradiance, since processed edges have uniform histograms. Because the valid range is known, we can determine to which image areas this partial calibration can be applied. In some cases, the complete image or entire object of interest could consist of intensities within the valid range, resulting in sufficient calibration.

Besides the appearance of two regions in a pixel area, defocus and motion blur are other possible causes of irradiance blending at edges. For a given edge region, as long as the point spread function of such blur encompasses at most two different scene regions of uniform intensity, the blending will remain a linear combination of irradiances as assumed in our technique. The acceptance or rejection of a window with blur is resolved by the selection criteria described in Section 3.

Some cameras may employ certain postprocessing operations, such as sharpening, which could alter the intensity distributions at edges. Since our edge selection technique attempts to exclude anomalous histograms, our method ideally will not locate any valid edge regions and consequently not calibrate such images. To better ensure the rejection of invalid edge regions, our technique for identifying incongruous histograms could be enhanced by more precise rejection criteria, such as using the DoRF database to bound the space of valid intensity histograms. Another potential cue for discerning an outlier histogram is that its shape may not transform in concordance with other histograms that overlap it in grayscale range. This remains an area for future investigation.

With the proposed statistical histogram measure for evaluating inverse response functions, radiometric calibration of grayscale images has been made possible in numerous instances where the camera is unknown. Although the method for single color images [7] is not applicable to grayscale images, our histogram approach may conversely be used in individual RGB channels of color images. Determining how histogram information may enhance color-based radiometric calibration is another interesting direction that we plan to examine.

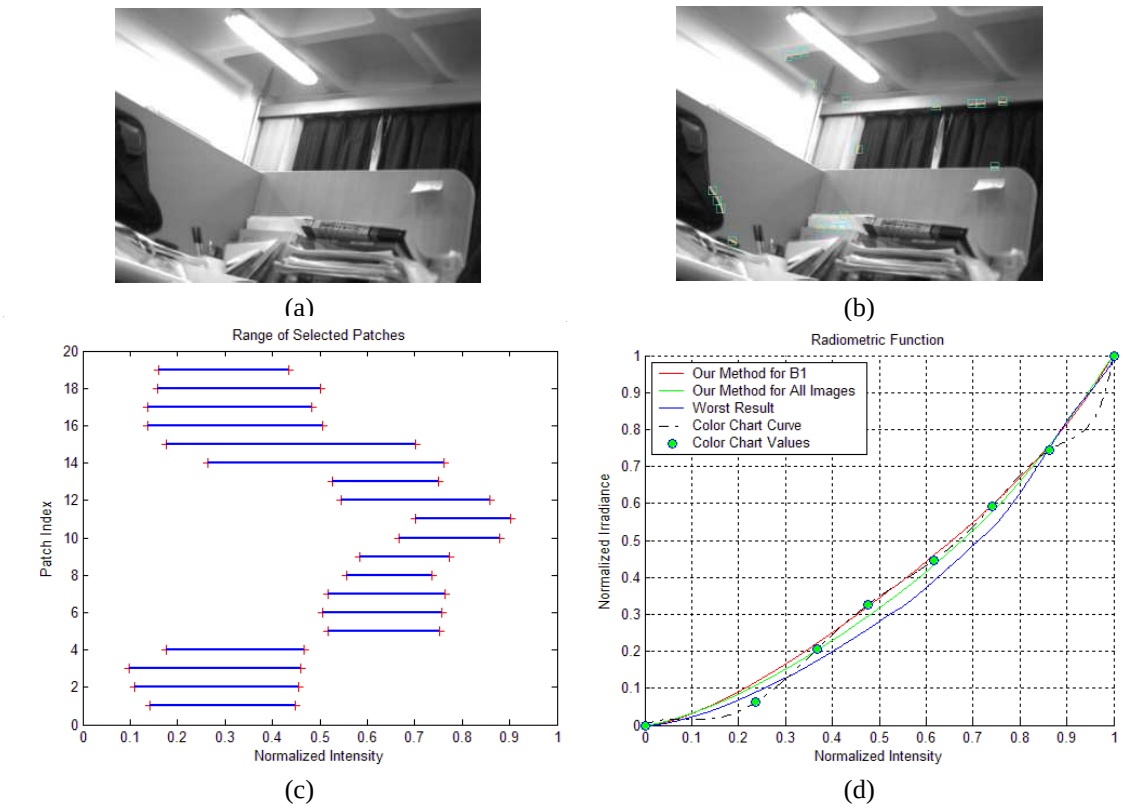

Figure 7. Results for a representative test image from an Ikegata IK-110 grayscale video camera. (a) Test image; (b) Detected edge patches; (c) Grayscale ranges of edge regions; (d) Inverse response functions.

Acknowledgement

The authors would like to thank Kyros Kutulakos for helpful discussions and Jinwei Gu for providing some useful code.

References

[1] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proc. ACM SIGGRAPH*, pages 369–378, 1997.

[2] H. Farid. Blind inverse gamma correction. *IEEE Trans. Image Processing*, 10(10):1428–1433, October 2001.

[3] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley, 1996.

[4] M. D. Grossberg and S. K. Nayar. Determining the camera response from images: What is knowable? *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(11):1455–1467, November 2003.

[5] M. D. Grossberg and S. K. Nayar. What is the space of camera response functions? In *Proc. IEEE Comp. Vision and Pattern Recog.*, pages II:602–609, 2003.

[6] S. J. Kim and M. Pollefeys. Radiometric alignment of image sequences. In *Proc. IEEE Comp. Vision and Pattern Recog.*, pages 645–651, 2004.

[7] S. Lin, J. Gu, S. Yamazaki, and H.-Y. Shum. Radiometric calibration from a single image. In *Proc. IEEE Comp. Vision and Pattern Recog.*, pages 938–945, 2004.

[8] S. Mann. Comparametric imaging: Estimating both the unknown response and the unknown set of exposures in a plurality of differently exposed images. In *Proc. IEEE Comp. Vision and Pattern Recog.*, pages I:842–849, 2001.

[9] S. Mann and R. Picard. Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. In *Proc. of IS&T, 48th annual conference*, pages 422–428, 1995.

[10] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Proc. IEEE Comp. Vision and Pattern Recog.*, pages II:374–380, 1999.

[11] S. K. Nayar and T. Mitsunaga. High dynamic range imaging: Spatially varying pixel exposures. In *Proc. IEEE Comp. Vision and Pattern Recog.*, pages I:472–479, 2000.

[12] C. Pal, R. Szeliski, M. Uyttendale, and N. Jojic. Probability models for high dynamic range imaging. In *Proc. IEEE Comp. Vision and Pattern Recog.*, pages 173–180, 2004.

[13] Y. Rubner, C. Tomasi, and L. J. Guibas. A metric for distributions with applications to image databases. In *Proc. Int. Conf. on Comp. Vision*, pages 59–66, 1998.

[14] Y. Tsin, V. Ramesh, and T. Kanade. Statistical calibration of ccd imaging process. In *Proc. Int. Conf. on Comp. Vision*, pages I:480–487, 2001.