

# Radiometric Self Calibration \*

Tomoo Mitsunaga  
Media Processing Laboratories  
Sony Corporation  
mitsunag@av.crl.sony.co.jp

Shree K. Nayar  
Department of Computer Science  
Columbia University  
nayar@cs.columbia.edu

## Abstract

A simple algorithm is described that computes the radiometric response function of an imaging system, from images of an arbitrary scene taken using different exposures. The exposure is varied by changing either the aperture setting or the shutter speed. The algorithm does not require precise estimates of the exposures used. Rough estimates of the ratios of the exposures (e.g. F-number settings on an inexpensive lens) are sufficient for accurate recovery of the response function as well as the actual exposure ratios. The computed response function is used to fuse the multiple images into a single high dynamic range radiance image. Robustness is tested using a variety of scenes and cameras as well as noisy synthetic images generated using 100 randomly selected response curves. Automatic rejection of image areas that have large vignetting effects or temporal scene variations make the algorithm applicable to not just photographic but also video cameras. Code for the algorithm and several results are publicly available at <http://www.cs.columbia.edu/CAVE/>.

## 1. Radiometric Response Function

Brightness images are inputs to virtually all computer vision systems. Most vision algorithms either explicitly or indirectly assume that brightness measured by the imaging system is linearly related to scene radiance. In practice, this is seldom the case. Almost always, there exists a non-linear mapping between scene radiance and measured brightness. We will refer to this mapping as the *radiometric response function* of the imaging system. The goal of this paper is to present a convenient technique for estimating the response function.

Let us begin by exploring the factors that influence the response function. In a typical image formation system, image irradiance  $E$  is related to scene radiance  $L$  as [8]:

$$E = L \frac{\pi}{4} \left(\frac{d}{h}\right)^2 \cos^4 \phi, \quad (1)$$

where,  $h$  is the focal length of the imaging lens,  $d$  is the diameter of its aperture and  $\phi$  is the angle subtended

by the principal ray from the optical axis. If our imaging system were ideal, the brightness it records would be  $I = E t$ , where,  $t$  is the time the image detector is exposed to the scene. This ideal system would have a linear radiometric response:

$$I = L k e, \quad (2)$$

where,  $k = \cos^4 \phi / h^2$  and  $e = (\pi d^2 / 4) t$ . We will refer to  $e$  as the *exposure* of the image, which could be altered by varying either the *aperture size*  $d$  or the *duration of exposure*  $t$ .

An image detector such as a CCD is designed to produce electrical signals that are linearly related to  $I$  [14]. Unfortunately, there are many other stages to image acquisition that introduce non-linearities. Video cameras often include some form of “gamma” mapping (see [13]). Further, the image digitizer, inclusive of A/D conversion and possible remappings, could introduce its own non-linearities. In the case of film photography, the film itself is designed to have a non-linear response [9]. Such film-like responses are often built into off-the-shelf digital cameras. Finally, the development of film into a slide or a print and its scanning into a digital format typically introduces further non-linearities (see [5]).

The individual sources of radiometric non-linearities are not of particular relevance here. What we do know is that the final brightness measurement  $M$  produced by an imaging system is related to the (scaled) scene radiance  $I$  via a response function  $g$ , i.e.  $M = g(I)$ . To map all measurements  $M$  to scaled radiance values  $I$ , we need to find the inverse function  $f = g^{-1}$ , where  $I = f(M)$ . Recovery of  $f$  is therefore the radiometric calibration problem. It is common practice to estimate  $f$  by showing the imaging system a uniformly lit calibration chart, such as the Macbeth chart [4], which includes patches of known relative reflectances. The known relative radiances  $I_i$  of the patches and the corresponding measurements  $M_i$  are samples that are interpolated to obtain an approximation to  $f$ . It is our objective to avoid such a calibration procedure which is only suited to controlled environment.

## 2. Calibration Without Charts

Recent work by Mann and Picard [11] and Debevec and Malik [5] has demonstrated how the radiometric response function can be estimated using images of ar-

\*This work was supported in parts by an ONR/DARPA MURI grant under ONR contract No. N00014-97-1-0553 and a David and Lucile Packard Fellowship. Tomoo Mitsunaga is supported by the Sony Corporation.

bitrary scenes taken under different known exposures. While the measured brightness values change with exposure, scene radiance values  $L$  remain constant. This observation permits the estimation of the inverse function  $f$  without prior knowledge of scene radiance. Once  $f$  has been determined, the images taken under different exposures can be fused into a single high dynamic range radiance image<sup>1</sup>.

Mann and Picard [11] use the function  $M = \alpha + \beta I^\gamma$  to model the response curve  $g$ . The bias parameter  $\alpha$  is estimated by taking an image with the lens covered, and the scale factor  $\beta$  is set to an arbitrary value. Consider two images taken under different exposures with known ratio  $R = e_1/e_2$ . The measurement  $M(I)$  in the first image ( $I$  is unknown) produces the measurement  $M(RI)$  in the second. A pixel with brightness  $M(RI)$  is sought in the first image that would then produce the brightness  $M(R^2I)$  in the second image. This search process is repeated to obtain the measurement series  $M(I)$ ,  $M(RI)$ , ...,  $M(R^nI)$ . Regression is applied to these samples to estimate the parameter  $\gamma$ . Since the model used by Mann and Picard is highly restrictive, the best one can hope for in the case of a general imaging system is a qualitative calibration result. Nonetheless, the work of Mann and Picard is noteworthy in that it brings to light the main challenges of this general approach to radiometric calibration.

Debevec and Malik [5] have developed an algorithm that can be expected to yield more accurate results. They use a sequence of high quality digital photographs (about 11) taken using precisely known exposures. This algorithm is novel in that it does not assume a restrictive model for the response function, and only requires that the function be smooth. If  $p$  and  $q$  denote pixel and exposure indices, we have  $M_{p,q} = g(E_p \cdot t_q)$  or  $\ln f(M_{p,q}) = \ln E_p + \ln t_q$ . With  $M_{p,q}$  and  $t_q$  known, the algorithm seeks to find discrete sample of  $\ln f(M_{p,q})$  and scaled irradiances  $E_p$ . These results are used to compute a high dynamic range radiance map. The algorithm of Debevec and Malik is well-suited when the images themselves are not noisy and precise exposure times are available.

Though our goals are similar to those of the above investigators, our approach is rather different. We use a flexible parametric model which can accurately represent a wide class of response functions encountered in practice. On the other hand, the finite parameters of the model allow us to recover the response function *without* precise exposure inputs. As we show, rough estimates (such as F-number readings on a low-quality lens) are sufficient to accurately calibrate the system as well as recover the actual exposure ratios. To ensure that our algorithm can

<sup>1</sup>For imaging systems with known response functions, the problem of fusing multiple images taken under different exposures has been addressed by several others (see [3], [10], [7] for examples). In [2], a CMOS imaging chip is described where each pixel measures the exposure time required to attain a fixed charge accumulation level. These exposure times can be mapped to a high dynamic range image.




Figure 1: Response functions of a few popular films and video cameras provided by their manufacturers. These examples illustrate that a high-order polynomial may be used to model the response function.

be applied to video systems, where it is easier to vary the aperture setting than the exposure time, we have implemented a pre-processing algorithm that rejects measurements with large vignetting effects and temporal changes during image acquisition. The recovered response function is used to fuse the multiple images into a high quality radiance map. We have conducted extensive testing of our algorithm by using noisy synthetic data. In addition, we have compared our calibration results for real images with those produced using calibration charts.

### 3. A Flexible Radiometric Model

It is worth noting that though the response curve can vary appreciably from one imaging system to the next, it is not expected to have exotic forms. The primary reason is that the detector output (be it film or solid state) is monotonic or at least semi-monotonic. Unless unusual mappings are intentionally built into the remaining components of the imaging system, measured brightness either increases or stays constant with increase in scene radiance (or exposure). This is illustrated by the response functions of a few popular systems shown in Figure 1. Hence, we claim that virtually any response function can be modeled using a high-order polynomial:

$$I = f(M) = \sum_{n=0}^N c_n M^n. \quad (3)$$

The minimum order of the polynomial required clearly depends on the response function itself. Hence, calibration could be viewed as determining the order  $N$  in addition to the coefficients  $c_n$ .

### 4. The Self-Calibration Algorithm

Consider two images of a scene taken using two different exposures,  $e_q$  and  $e_{q+1}$ , where  $R_{q,q+1} = e_q/e_{q+1}$ . The ratio of the scaled radiance at any given pixel  $p$  can

written using expression (2) as:

$$\frac{I_{p,q}}{I_{p,q+1}} = \frac{L_p k_p e_q}{L_p k_p e_{q+1}} = R_{q,q+1}. \quad (4)$$

Hence, the response function of the imaging system is related to the exposure ratio as:

$$\frac{f(M_{p,q})}{f(M_{p,q+1})} = R_{q,q+1}. \quad (5)$$

We order our images such that  $e_q < e_{q+1}$  and hence  $0 < R_{q,q+1} < 1$ . Substituting our polynomial model for the response function, we have:

$$\frac{\sum_{n=0}^N c_n M_{p,q}^n}{\sum_{n=0}^N c_n M_{p,q+1}^n} = R_{q,q+1}. \quad (6)$$

The above relation may be viewed as the basis for the joint recovery of the response function and the exposure ratio. However, an interesting ambiguity surfaces at this point. Note that from (5) we also have  $(f(M_{p,q})/f(M_{p,q+1}))^u = R_{q,q+1}^u$ . This implies that, in general,  $f$  and  $R$  can only be recovered up to an unknown exponent  $u$ . In other words, an infinite number of  $f$ - $R$  pairs would satisfy equation (5).

Interestingly, this  $u$ -ambiguity is greatly alleviated by the use of the polynomial model. Note that, if  $f$  is a polynomial,  $f^u$  can be a polynomial only if  $u = v$  or  $u = 1/v$  where  $v$  is a natural number, i.e.  $u = (., 1/3, 1/2, 1, 2, 3, .)$ . This by no means implies that, for any given polynomial, all these multiple solutions must exist. For instance, if  $f(M) = M^3$ ,  $u = 1/2$  does not yield a polynomial. On the other hand,  $u = 1/3$  does result in  $f(M) = M$  which, in turn, can have its own  $u$ -ambiguities. In any case, the multiple solutions that arise are well-spaced with respect to each other. Shortly, the benefits of this restriction will become clear.

For now, let us assume that the exposure ratios  $R_{q,q+1}$  are known to us. Then, the response function can be recovered by formulating an error function that is the sum of the squares of the errors in expression (6):

$$\mathcal{E} = \sum_{q=1}^{Q-1} \sum_{p=1}^P \left[ \sum_{n=0}^N c_n M_{p,q}^n - R_{q,q+1} \sum_{n=0}^N c_n M_{p,q+1}^n \right]^2, \quad (7)$$

where,  $Q$  is the total number of images used. If we normalize all measurements such that  $0 \leq M \leq 1$  and fix the indeterminable scale using  $f(1) = I_{max}$ , we get the additional constraint:

$$c_N = I_{max} - \sum_{n=0}^{N-1} c_n. \quad (8)$$

The response function coefficients are determined by solving the system of linear equations that result from setting:

$$\frac{\partial \mathcal{E}}{\partial c_n} = 0. \quad (9)$$

In most inexpensive imaging systems, photographic or video, it is difficult to obtain accurate estimates of the exposure ratios  $R_{q,q+1}$ . The user only has access to the F-number of the imaging lens or the speed of the shutter. In consumer products these readings can only be taken as approximations to the actual values. In such cases, the restricted  $u$ -ambiguity provided by the polynomial model proves valuable. Again, consider the case of two images. If the initial estimate for the ratio provided by the user is a reasonable guess, the actual ratio is easily determined by searching in the vicinity of the initial estimate; we search for the  $R$  that produces the  $c_n$  that minimize  $\mathcal{E}$ . Since the solution for  $c_n$  is linear, this search is very efficient.

However, when more than two images are used the dimensionality of the search for  $R_{q,q+1}$  is  $Q - 1$ . When  $Q$  is large, the search can be time consuming. For such cases, we use an efficient iterative scheme where the current ratio estimates  $R_{q,q+1}^{(k-1)}$  are used to compute the next set of coefficients  $c_n^{(k)}$ . These coefficients are then used to update the ratio estimates using (6):

$$R_{q,q+1}^{(k)} = \sum_{p=1}^P \frac{\sum_{n=0}^N c_n^{(k)} M_{p,q}^n}{\sum_{n=0}^N c_n^{(k)} M_{p,q+1}^n}, \quad (10)$$

where, the initial ratio estimates  $R_{q,q+1}^{(0)}$  are provided by the user. The algorithm is deemed to have converged when:

$$|f^{(k)}(M) - f^{(k-1)}(M)| < \epsilon, \quad \forall M, \quad (11)$$

where  $\epsilon$  is a small number.

It is hard to embed the recovery of the order  $N$  into the above algorithm in an elegant manner. Our approach is to place an upper bound on  $N$  and run the algorithm repeatedly to find the  $N$  that gives the lowest error  $\mathcal{E}$ . In our experiments, we have used an upper bound of  $N=10$ .

## 5. Evaluation: Noisy Synthetic Data

Shortly, we will present experiments with real scenes that demonstrate the performance of our calibration algorithm. However, a detailed analysis of the behavior of the algorithm requires the use of noisy synthetic data. For this, 100 monotonic response functions were generated using random numbers for the coefficients  $c_n$  of a fifth-order polynomial. Using each response function, four synthetic images were generated with random radiance values and random exposure ratios in the range  $0.45 \leq R \leq 0.55$ . This range of ratio values is of particular interest to us because in almost all commercial lenses and cameras a single step in the F-number setting or shutter speed setting results in an exposure ratio of approximately 0.5. The pixel values were normalized such that  $0 \leq M \leq 1$ . Next, normally distributed noise with  $\sigma = 0.005$  was added to each pixel value. This translates to  $\sigma = 1.275$  gray levels when  $0 \leq M \leq 255$ , as in the case of a typical imaging system. The images were then




Figure 2: Self-calibration was tested using 100 randomly generated response functions. Here a few of the recovered (solid) and actual (dots) response functions are shown. In each case, four noisy test images were generated using random exposure ratios between image pairs in the range  $0.45 \leq R \leq 0.55$ . The initial ratio estimates were chosen to be 0.5.




Figure 3: The percentage average error between actual and estimated response curves for the 100 synthetic trials. The maximum error was found to be 2.7%.

quantized to have 256 discrete gray levels. The algorithm was applied to each set of four images using initial exposure ratios of  $R_{q,q+1}^{(0)} = 0.5$ .

All 100 response functions and exposure ratios were accurately estimated. Figure 2 shows a few of the actual (dots) and computed (solid) response functions. As can be seen, the algorithm is able to recover a large variety of response functions. In Figure 3 the percentage of the average error (over all  $M$  values) between the actual and estimated curves are shown for the 100 cases. Despite the presence of noise, the worst-case error was found to be less than 2.7%.

Figure 4 shows the error  $\mathcal{E}$  plotted as a function of exposure ratio  $R$ . In this example, only two images are used and the actual and initial exposure ratios are 0.7 and 0.625, respectively. The multiple local minima correspond to solutions that arise from the  $u$ -ambiguity described before. The figure illustrates how the ratio value converges from the initial estimate to the final one. In all our experiments, the algorithm converged to the correct solution in less than 10 iterations.

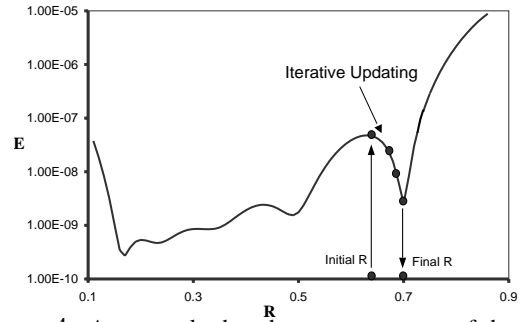


Figure 4: An example that shows convergence of the algorithm to the actual exposure ratio (0.7) from a rough initial ratio estimate (0.625), in the presence of the  $u$ -ambiguity.

## 6. Implementation

We now describe a few issues that need to be addressed while implementing radiometric self-calibration.

### 6.1. Reducing Video Noise

Particularly in the context of video, it is important to ensure that the data that is provided to the self-calibration algorithm has minimal noise. To this end, we have implemented a pre-processing step that uses temporal and spatial averaging to obtain robust pixel measurements. Noise arises from three sources, namely, electrical read-out from the camera, quantization by the digitizer hardware [6], and motion of scene objects during data acquisition. The random component of the former two noise sources can be reduced by temporal averaging of  $t$  images (typically  $t = 100$ ). The third source can be omitted by selecting pixels of spatially flat area. To check spatial flatness, we assume normally distributed noise  $N(0, \sigma^2)$ . Under this assumption, if an area is spatially flat,  $sS^2/\sigma^2$  can be modeled by the  $\chi^2$  distribution [12], where  $S$  is the spatial variance of  $s$  pixel values (typically  $s = 5 \times 5$ ). We approximate  $S^2$  with the spatial variance  $\sigma_s^2$  of the temporally averaged pixel values, and  $\sigma^2$  with the temporal variance  $\sigma_t^2$ . Therefore, only those pixels are selected that pass the following test:

$$\frac{s \sigma_s^2}{\sigma_t^2} \leq \chi^2_{s-1}(\psi), \quad (12)$$

where,  $\psi$  is the rejection ratio which is set to 0.05. When temporal averaging is not applied,  $t=1$  and  $\sigma_t$  is set to 0.01.

### 6.2. Vignetting

In most compound lenses vignetting increases with the aperture size [1]. This introduces errors in the measurements which are assumed to be only effected by exposure changes. Vignetting effects are minimal at the center of the image and increase towards the periphery. We have implemented an algorithm that robustly detects pixels that are corrupted by vignetting. Consider two consecutive images  $q$  and  $q + 1$ . Corresponding brightness measurements  $M_{p,q}$  and  $M_{p,q+1}$  are plotted against each other. In the absence of vignetting, all pixels with the same measurement value in one image should produce




Figure 5: Fusing multiple images taken under different exposures into a single scaled radiance image.

equal measurement values in the second image, irrespective of their locations in the image. The vignetting-free area for each image pair is determined by finding the smallest image circle within which the  $M_{p,q}$ - $M_{p,q+1}$  plot is a compact curve with negligible scatter.

### 6.3. High Dynamic Range Radiance Images

Once the response function of the imaging system has been computed, the  $Q$  images can be fused into a single high dynamic range radiance image, as suggested in [11] and [5]. This procedure is simple and is illustrated in Figure 5. Three steps are involved. First, each measurement  $M_{p,q}$  is mapped to its scaled radiance value  $I_{p,q}$  using the computed response function  $f$ . Next, the scaled radiance is normalized by the scaled exposure  $\tilde{e}_q$  so that all radiance value end up with the same effective exposure. Since the absolute value of each  $\tilde{e}_q$  is not important for normalization, we simply compute the scaled exposures so that their arithmetic mean is equal to 1. The final radiance value at a pixel is then computed as a weighted average of its individual normalized radiance values. In previous work, a hat function [5] and the gradient of the response function [11] were used for the weighting function. Both these choices are somewhat ad-hoc, the latter less so than the former. Note that a measurement can be trusted most when its signal-to-noise ratio (SNR) as well as its sensitivity to radiance changes are maximum. The SNR for the scaled radiance value  $I$  is:

$$\text{SNR} = I \frac{dM}{dI} \frac{1}{\sigma_N(M)} = \frac{f(M)}{\sigma_N(M) f'(M)}, \quad (13)$$

where,  $\sigma_N(M)$  is the standard deviation of the measurement noise. Using the assumption that the noise  $\sigma_N$  in (13) is independent of the measurement pixel value  $M$ , we can define the weighting function as:

$$w(M) = f(M)/f'(M). \quad (14)$$

### 6.4. Handling Color

In the case of color images, a separate response function is computed for each color channel (red, green and blue). This is because, in principle, each color channel could have its own response to irradiance. Since each of the response functions can only be determined up to a scale factor, the relative scalings between the three computed radiance images remain unknown. We resolve this problem by assuming that the three response

functions preserve the chromaticity of scene points. Let the measurements be denoted as  $\mathbf{M} = [M_r, M_g, M_b]^T$  and the computed scaled radiance values be  $\mathbf{I} = [I_r, I_g, I_b]^T$ . Then, the color-corrected radiance values are  $\mathbf{I}_c = [k_r I_r, k_g I_g, k_b I_b]^T$  where,  $k_r/k_b$  and  $k_g/k_b$  are determined by applying least-squares minimization to the chromaticity constraint  $\mathbf{I}_c / \|\mathbf{I}_c\| = \mathbf{M} / \|\mathbf{M}\|$ .

## 7. Experimental Results

The source code for our self-calibration algorithm and several experimental results are made available at <http://www.cs.columbia.edu/CAVE/>. Here, due to limited space, we will present just a couple of examples.

Figure 6 shows results obtained using a Canon Optura video camera. In this example, the gray scale output of the camera was used. Figure 6(a) shows 4 of the 5 images of an outdoor scene obtained using five different F-number settings. As can be seen, the low exposure images are too dark in areas of low scene radiance and the high exposure ones are saturated in bright scene regions. To reduce noise, temporal averaging using  $t=100$  and spatial averaging using  $s=3 \times 3$  were applied. Next, automatic vignetting detection was applied to the 5 averaged images. Figure 6(b) shows the final selected pixels (in black) for one of the image pairs.

The solid line in Figure 6(d) is the response function computed from the 5 scene images. The initial exposure ratio estimates were set to 0.5, 0.25, 0.5 and 0.5 and the final computed ratios were 0.419, 0.292, 0.570 and 0.512. These results were verified using a Macbeth calibration chart with patches of known relative reflectances (see Figure 6(c)). The chart calibration results are shown as dots which are in strong agreement with the computed response function. Figure 6(e) shows the radiance image computed using the 5 input images and the response function. Since conventional printers and displays have dynamic ranges of 8 or less bits, we have applied histogram equalization to the radiance image to bring out the details within it. The small windows on the sides of the radiance image show further details; histogram equalization was applied locally within each window.

Figure 7 shows two results obtained using a Nikon 2020 film camera. In each case 5 images were captured of which only 4 are shown in Figures 7(a) and 7(d). In the first experiment, Kodachrome slide film with ISO 200 was used. These slides were taken using F-number = 8 and approximate (manually selected) exposure times of 1/30, 1/15, 1/8, 1/2 and 1 (seconds). In the second experiment, the pictures were taken with the same camera but using F-number = 11 and ISO 100 slide film. The exposure settings were 1/500, 1/250, 1/125, 1/60 and 1/30. The developed slides were scanned using a Nikon LS-3510AF scanner that produces a 24 bit image (8 bits per color channel).

The self-calibration algorithm was applied separately to each color channel (R, G, B), using initial ratio estimates of 0.5, 0.25, 0.5 and 0.5 in the first experiment and all

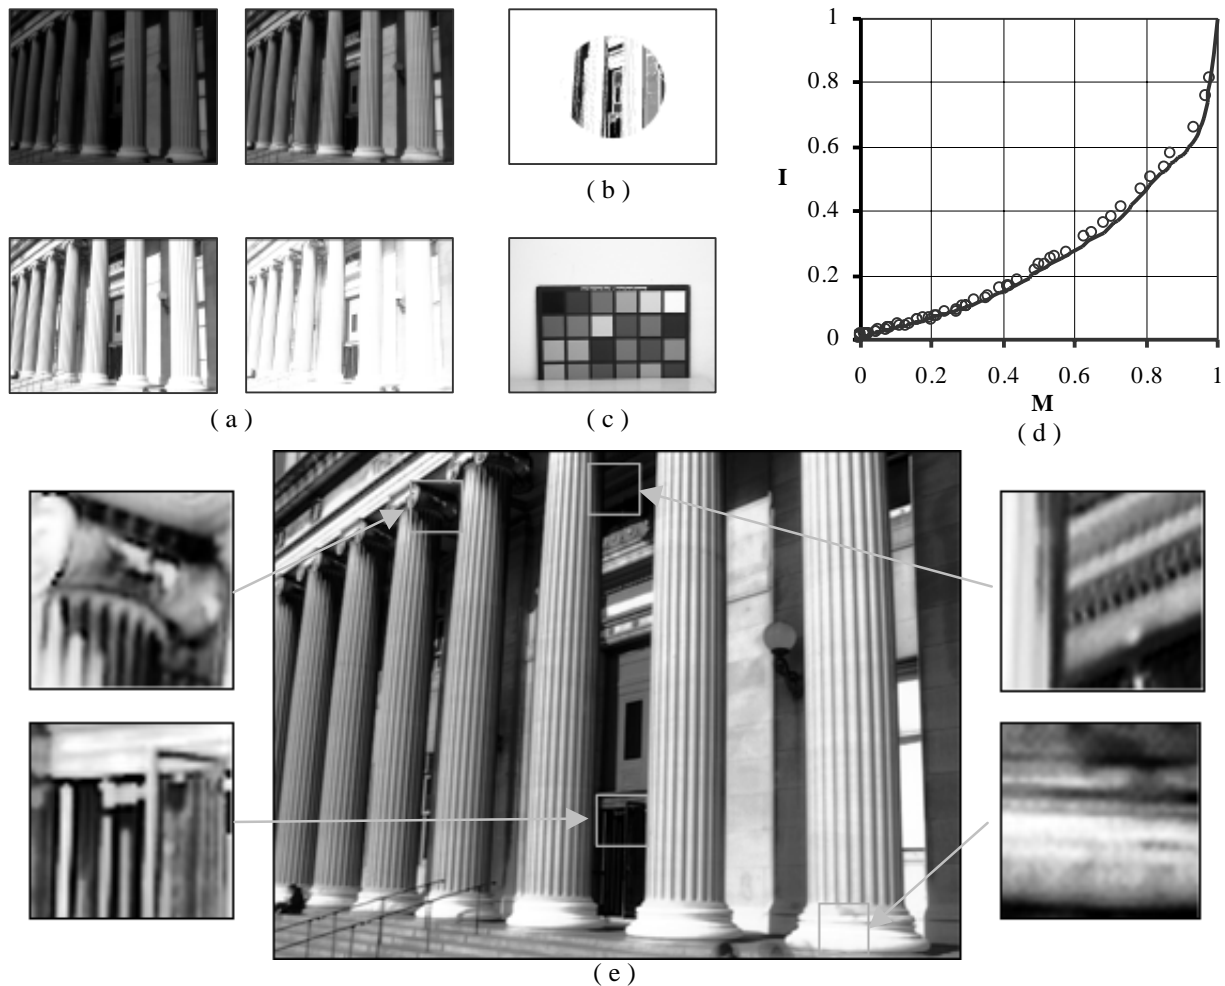


Figure 6: (a) Self-calibration results for gray scale video images taken using a Canon Optura camera. (b) Temporal averaging, spatial averaging and vignetting detection are used to locate pixels (shown in black) that produce robust measurements. (c) The self-calibration results are verified using a uniformly lit Macbeth color chart with patches of known reflectances. (d) The computed response function (solid line) is in strong agreement with the chart calibration results (dots). (e) The computed radiance image is histogram equalized to convey some of the details it includes. The image windows on the two sides of the radiance image are locally histogram equalized to bring forth further details.

ratios set to 0.5 in the second. The three (R, G, B) computed response functions are shown in Figures 7(b) and 7(e). A radiance image was computed for each color channel and then the three images were scaled to preserve chromaticity, as described in section 6.4.. The final radiance images shown in Figure 7(c) and 7(f) are histogram equalized. As before, the small windows within the radiance images show further details brought out by local histogram equalization. Noteworthy are the details of the lamp and the chain on the wooden window in Figure 7(c) and the clouds in the sky and the logs of wood inside the clay oven in Figure 7(f).

For the exploration of high dynamic range images, we have developed a simple interactive visualization tool, referred to as the “Detail Brush.” It is a window of adjustable size and magnification that can be slid around the image while histogram equalization within the window is performed in real time. This tool is also available

at <http://www.cs.columbia.edu/CAVE/>.

## References

- [1] N. Asada, A. Amano, and M. Baba. Photometric Calibration of Zoom Lens Systems. *Proc. of IEEE International Conference on Pattern Recognition (ICPR)*, pages 186–190, 1996.
- [2] V. Brajovic and T. Kanade. A Sorting Image Sensor: An Example of Massively Parallel Intensity-to-Time Processing for Low-Latency Computational Sensors. *Proc. of IEEE Conference on Robotics and Automation*, pages 1638–1643, April 1996.
- [3] P. Burt and R. J. Kolczynski. Enhanced Image Capture Through Fusion. *Proc. of International Conference on Computer Vision (ICCV)*, pages 173–182, 1993.

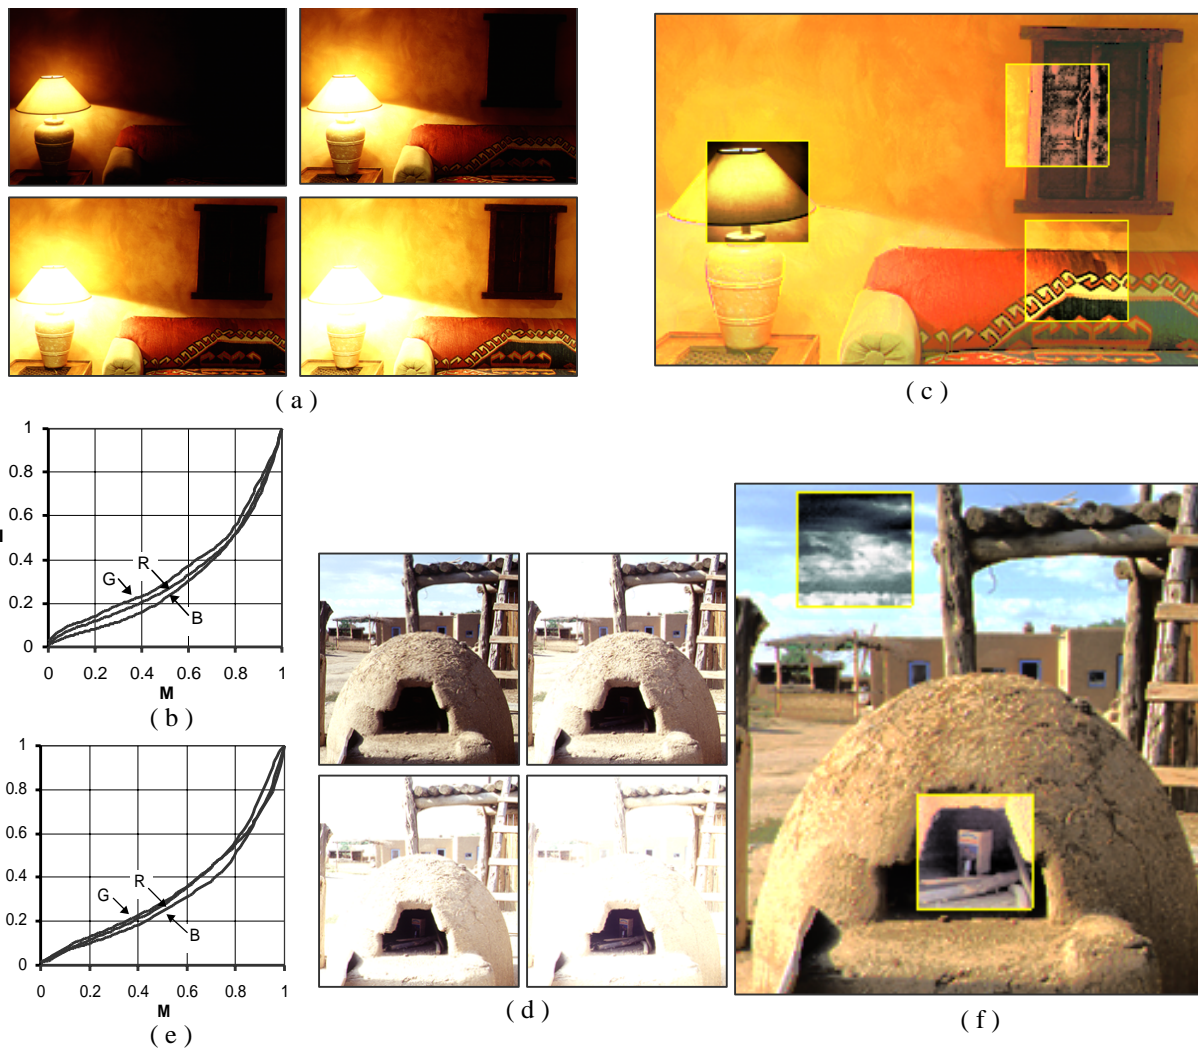


Figure 7: Self-calibration results for color slides taken using a Nikon 2020 photographic camera. The slides are taken using film with ISO rating of 200 in (a) and 100 in (d). Response functions are computed for each color channel and are shown in (b) and (e). The radiance images in (c) and (f) are histogram equalized and the three small windows shown in each radiance image are locally histogram equalized. The details around the lamp and within the wooden window in (c) and the clouds in the sky and the objects within the clay oven in (f) illustrate the richness of visual information embedded in the computed radiance images.

- [4] Y.-C. Chang and J. F. Reid. RGB Calibration for Analysis in Machine Vision. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 5(10):1414–1422, October 1996.
- [5] P. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. *Proc. of ACM SIGGRAPH 1997*, pages 369–378, 1997.
- [6] G. Healey and R. Kondepudy. Radiometric CCD Camera Calibration and Noise Estimation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 16(3):267–276, March 1994.
- [7] K. Hirano, O. Sano, and J. Miyamichi. A Detail Enhancement Method by Merging Multi Exposure Shift Images. *IEICE*, J81-D-II(9):2097–2103, 1998.
- [8] B. K. P. Horn. *Robot Vision*. MIT Press, Cambridge, MA, 1986.
- [9] T. James, editor. *The Theory of the Photographic Process*. Macmillan, New York, 1977.
- [10] B. Madden. Extended Intensity Range Imaging. Technical Report MS-CIS-93-96, Grasp Laboratory, University of Pennsylvania, 1996.
- [11] S. Mann and R. Picard. Being ‘Undigital’ with Digital Cameras: Extending Dynamic Range by Combining Differently Exposed Pictures. *Proc. of IST’s 48th Annual Conference*, pages 422–428, May 1995.
- [12] S. L. Meyer. *Data Analysis for Scientists and Engineers*. John Wiley and Sons, 1992.
- [13] C. A. Poynton. *A Technical Introduction to Digital Video*. John Wiley and Sons, 1996.
- [14] A. J. P. Theuwissen. *Solid State Imaging with Charge-Coupled Devices*. Kluwer Academic Press, Boston, 1995.