

Exposure time estimation for high dynamic range imaging with hand held camera

Lukáš Cerman and Václav Hlaváč

Czech Technical University, Faculty of Electrical Engineering, Center for Machine Perception
121 35 Prague 2, Karlovo náměstí 13, Czech Republic
cermal1@fel.cvut.cz

Abstract *The limited dynamic range of a camera may be extended by composing differently exposed images of the same scene. The nonlinear camera has to be calibrated radiometrically first. We show that the calibration process can be difficult for real cameras. The improvement can be achieved if linear 12-bit RAW images are used as offered by modern mid-class and professional cameras. The knowledge of exposure time positively affects the radiometric quality of the composed high dynamic range (HDR) image. This knowledge also helps in registration of differently exposed hand held shots. This contribution presents a new method for estimating exposure times from the histograms of images captured using a linear response camera or generated from a RAW image. The presented method does not require spatially aligned images. The actual process of HDR image composition needs perfectly aligned images on its input. We present a method for registering differently exposed images captured with a hand held camera. The presented registration method enables capture of HDR images without the need of mounting a camera on the tripod. The methods are implemented in Matlab and tested on real data.*

1 Motivation

Since the early years of photography, the photographic process has been limited by the dynamic range of a photosensitive material which is able to record only a limited range of light intensity.

Imagine a high contrast scene containing dark shadows in some parts and very bright zones in the other parts of the scene. If someone sets the exposure with the aim to capture details in the dark zones correctly, like in Figure 1a, then details in the bright zones are lost by overexposure. The material is saturated and is not able to record differences between light intensities above the saturation level. Likewise, if someone exposes to capture bright zones correctly, like in Figure 1b, then details in the underexposed dark zones are lost in sensor noise or a film grain.

Many techniques for setting proper exposure were developed during years, probably the most popular is the legendary Zone System for black and white photography by Ansel Adams [1]. This is a sophisticated method of understanding and controlling the exposure, development of the negative and the print process. Similarly, the modern automatic cameras have a very sophisticated metering and expo-




Figure 1: The same scene exposed for shadows (a) and lights (b).

sure programs developed by the analysis of thousands professional photographs. However, all such systems, no matter how sophisticated, do not increase the limited dynamic range of the sensing material. They are only trying to exploit it best by the correct exposure and image processing.

Obviously, the limited dynamic range is not a problem of artistic photography only, but affects industrial and scientific processes having images as their inputs. Such processes become more robust if they can utilize the whole dynamic range of a high contrast scene.

2 Problem Specification

Although the HDR image can be captured directly using special HDR sensors like Kodak KAC-9628 [11], IMS Chips HDRC sensors [10], Silicon Vision Products [13], SMaL Camera [14] and Pixim [12], these solutions are expensive. However, it can be shown that the HDR image can be obtained with widely available cameras by fusing multiple differently exposed images of the same scene.

2.1 Combining Differently Exposed Images

The idea of fusing images with different exposures is not new. In the early sixties, Charles Wyckoff developed a wide dynamic range film composed of three layers of the same spectral sensitivity but differing in their overall sensitivity to light. The three layers were colored exactly like in the color negative film – cyan, magenta and yellow – so after printing on the color paper Wyckoff created a pseudocolor HDR image where the individual gray¹ tones were represented by different colors.

The same principle may be applied using a standard film or an image sensor. Multiple images of the same scene differing only in exposure can be captured. Each captured image records different range of light intensities. If there is enough overlapping exposures ranging from dark images,

¹Note that all three layers have the same spectral sensitivity, so the recorded information is achromatic.




Figure 2: Scheme of a camera illustrating how the scene radiance is transformed to pixel values.

recording information in the bright zones, to light images, recording information in the dark zones, an HDR image can be reconstructed by properly combining the information from captured images. This combined image will represent the whole scene radiance in one floating point map.

2.2 Camera Response Function and HDR Fusion

There is a camera scheme in Figure 2. Let us explore what the camera measures. Generally, the pixel values are proportional to the scene radiance, transformed by some nonlinear mapping called the *response function*. This nonlinearity is induced by different parts of the imaging process.

As the light – characterized by scene radiance L – is passing the lens, it is attenuated by the lens aperture and yields sensor irradiance E . The total quantity of light accumulated on the image sensor during integration time Δt controlled by the shutter produces the sensor exposure

$$I = E\Delta t. \quad (1)$$

Typical sensors like CCD or CMOS are designed to produce electrical signals which are linearly proportional to the sensor exposure up to some saturation level. Above this level, the sensor is not able to distinguish between different exposure values.

In the case of film photography, the film itself is designed to have a nonlinear response and the development process causes further nonlinearities.

After exposure, the accumulated charge is converted to integer values using analog-digital converters connected to each photosensitive element on the image sensor. The process of digitization brings quantization noise to recorded data. In addition, the digital values are in the case of the typical camera further transformed to mimic film characteristics or scaled by a γ function to properly display on CRT monitors. At this point, we must note that many modern cameras offer a RAW image output. The RAW image does not suffer from in-camera software postprocessing. It is affected only by sensor saturation and quantization noise.

Individual sources of nonlinearity in the imaging process are not important and the whole process can be represented by one nonlinear function – a camera response function f . The measured pixel values M are proportional to the sensor exposure I scaled by the response function f , $M = f(I)$.

To fuse an HDR image from a set of differently exposed images, all source images have to be transformed to the same domain – irradiance E . There is an assumption that the scene radiance did not change between the individual exposures. To map pixel values M to irradiance values E one

needs to find inverse function $g = f^{-1}$ so that the sensor exposure $I = g(M)$. Knowing the integration time Δt irradiance E can be obtained using expression (1).²

2.3 Image Registration

Composition of an HDR image from a set of differently exposed images of the same scene requires a perfect alignment of the individual images. This is no problem when capturing from a tripod. However, for a hand held capturing there is a need to perform a dense registration of the source images – i.e. to find pixel-to-pixel correspondences between the individual exposures.

Unlike a usual registration problem, there is a need to register images with large exposure differences and small spatial difference. The spatial difference is only caused by a little camera shift and shakes between the individual shots. The knowledge of the camera response may help this process by transforming all differently exposed images in the same domain – irradiance E .

Hand held capturing is also limited by the need of sufficiently long exposure times preventing the individual images from vibration blur. This limitation may be removed by applying multichannel deconvolution to blurred images but it is out of the scope of this work.

3 Related Work

3.1 Camera Calibration

The camera response function can be estimated by taking a picture of the uniformly illuminated chart containing patches of known reflectance, such as the Gretag Macbeth chart [2]. However this process is quite complicated and can be performed only in laboratory conditions which are not always accessible. In addition, the camera response may be affected by temperature changes requiring frequent recalibration.

Fortunately, it has been shown that a set of differently exposed images contains usually enough information to recover a response using the images themselves [7], [3], [8], [15], [6] and [4].

Mann and Pickard [7] modeled g as a γ function. Debevec and Malik [3] used a nonparametric model of g , sampled at certain values, and represented it by a vector. Mitsunaga and Nayar [8] assumed that the response can be modeled by a low degree polynomial. Further, they are able to start from rough estimates of exposure ratios k and recover both g and k by an iterative approach.

The previous methods used pixel values selected from images at corresponding points to constrain the solution of g . Selection of these points had a significant impact on the quality of the solution. Mann et al. showed [6] that all information for response recovery available in the image pair can

² The use of radiance and irradiance here is not entirely correct in terms of radiometric quantity because the sensor does not have uniform spectral response. The sensor measures quantity of light $I(\lambda)$ integrated over the full range of wavelengths λ weighted by its spectral response $R(\lambda)$ – $\int_{\lambda} I(\lambda)R(\lambda)d\lambda$. However, a typical sensor will not have the same spectral response as the human eye, so the measured quantity cannot be luminance either. In this text, we will use terms radiance and irradiance, while reminding the reader that the quantities we are dealing with are weighted by the spectral response of the imaging system.

be represented by the *brightness transfer function* T . The function describes how the brightness transfers from one image to another, $M_B = T(M_A)$.

Grossberg and Nayar showed [4] that T can be recovered using image histograms only, requiring no registration. They modified the solution of Mitsunaga and Nayar [8] to constrain the polynomial model by values of T .

3.2 Image Registration

Techniques used to register differently exposed images before an HDR image composition are similar to those used in mosaic and panorama stitching applications. Schechner and Nayar [9] mentioned traditional registration methods based on optimizing the sum of square or absolute differences, correlation or mutual information between frames, over the general motion parameters.

Kang et al. [5] implemented a sophisticated registration algorithm which was used in their HDR video application. Before registration, they boosted the brightness of darker frames to bring all frames in the video stream to the same brightness level. To register exposure compensated frames, they used motion modeling together with hierarchical homography. The motion model is used to compensate camera movement while the hierarchical homography is used to compensate for camera shake.

Ward described [16] a fast image registration method for HDR image composition. The median thresholded images are registered using his approach. The registration process is performed using the fast bit manipulation routines applied on an image pyramid. However, to keep this method fast it is possible to compensate for image shifts only and not for rotation or general homography which would require image warping. For the same reason it is not possible to register shifts at subpixel precision.

4 Estimate of the Camera Response Function

4.1 Implemented Calibration Methods

For a comparison and a performance test, we decided to implement more than one response recovery method. We implemented both versions of polynomial approximation – the original one by Mitsunaga and Nayar [8] and the histogram based modification by Grossberg and Nayar [4]. We also implemented the original nonparametric method of Debevec and Malik [3] according to their sample implementation, which was published in the appendix of their article.

Additionally, we modified the original method of Debevec and Malik [3] to recover a nonparametric model of g from image histograms. We followed the idea of Grossberg and Nayar [4] originally used for the polynomial model of Mitsunaga and Nayar.

4.2 Experimental Verification

To verify the usability of response recovery methods described in articles [3], [8] and [4] we tested our implementation of these methods on synthetic images created with the known response function.

4.2.1 Synthetic Data Without Noise We created six virtual grayscale images generated from a green channel of an

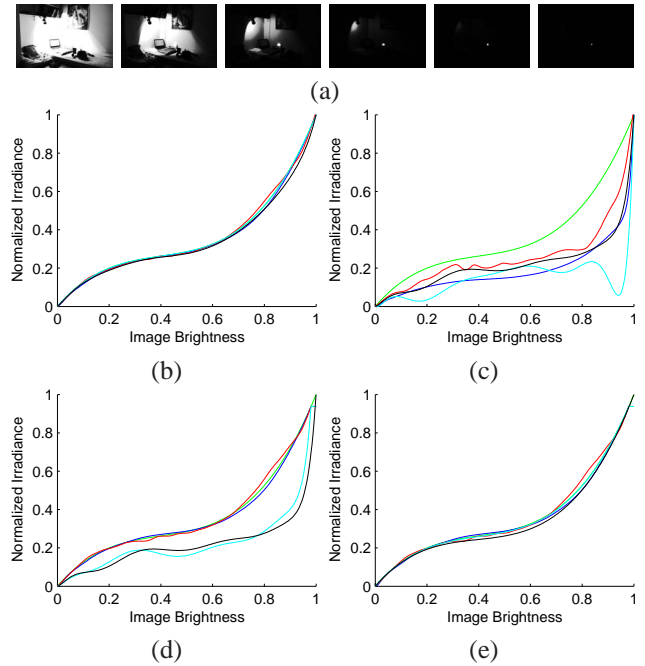


Figure 3: (a) Synthetic nonlinear images. Plots: Green line represents inverse of camera response function actually used. The other lines represent curves estimated by using different algorithms – nonparametric model of Debevec and Malik (blue), histogram based version of the same algorithm (red), polynomial model by Mitsunaga and Nayar (cyan) and histogram based version of this algorithm (black). (b) estimates from data without noise and quantization. (c)–(e) estimates from noisy and quantized data.

HDR image. The individual images were generated using formula $M_{ij} = g^{-1}(s(E_i \Delta t_j))$, where E_i is normalized irradiance value of pixel i . Δt_j is exposure time of image j , we used values $\{1, 4, 16, 64, 256, 1024\}$ which provide 2 stops difference between the individual exposures. $s(I)$ is saturation function

$$s(I) = \begin{cases} I & \text{for } I \leq 1 \\ 1 & \text{for } I > 1 \end{cases} \quad (2)$$

g^{-1} is the known camera response function, we used a 3rd order polynomial. M_{ij} is the generated brightness value of pixel i in the virtual image j .

The results of different response recovery algorithms are plotted in Figure 3b. The true response function g is painted by the green line. The original method by Debevec and Malik [3] is painted blue, the modified version of their algorithm utilizing only image histograms is displayed in red.

We used 10th order polynomial to estimate camera response using the polynomial model of Mitsunaga and Nayar [8]. The original method of Mitsunaga and Nayar is shown in cyan. The histogram based version of this method described by Grossberg and Nayar [4] is shown in black.

It is obvious that the results of all methods are very similar and close to the true response g when applied to data without noise and quantization.

4.2.2 Noisy Synthetic Data The previous experiment verified the ability of all tested algorithms to recover camera response, but only for images without noise and quantization errors. However, the real situation is rather different. To

make our experiment more natural, the images in the testing series were quantized to 256 discrete levels. A normally distributed noise with zero mean and unit variance was added to each pixel value. Noisy values smaller than 0 were clamped to 0. Similarly, values over 1 were clamped to 1.

We used the same algorithms to estimate the response function from the noisy data like in the previous experiment. The resulting estimates are plotted in Figure 3c. It can be seen that all methods used failed to recover the true camera response shown by the green line.

Duiker, Hawkins, and Debevec provided the program `mkhdr` based on the algorithm by Debevec and Malik [3]. The program works quite well even if quantization errors and noise are present. When we closely examined the response function generated by `mkhdr`, we learned that it is cut above some unknown level. The level was different for each color channel and varied between different image series. According to this new knowledge, we decided to recover response function only from values in the range $(0, 0.98)$. The value of response function at interval $(0.98, 1)$ was set to be constant value of $g(0.98)$. The results of such an improved algorithm were very similar to those of `mkhdr` utility. The results of this improved algorithm with noisy synthetic data from previous experiment are plotted in Figure 3d. We see that the two methods based on the non-parametric model of Debevec and Malik provided results that were very close to the real response function even for noisy and quantized data. However, the polynomial approximation still failed to recover the true camera response.

In the last experiment, we tried to estimate the response function using a polynomial of the same order (third) as the real response function. It can be seen in Figure 3e that the third order polynomial models fitted well to the real response function.

4.2.3 Summary of Response Recovery Test It is obvious from the experiments presented above that the naive implementation of the algorithms described in articles [3], [8] and [4] does not give the reasonable results. There are many unknown aspects affecting the quality of results provided by these algorithms.

We showed that the method of Debevec and Malik [3] can be strengthened by ignoring the brightness values greater than some value near the saturation point. Additionally, we showed that the polynomial based methods are, in practise, useful to model polynomial response of a known order only. Is question, how the estimate improves if the monotocity was enforced to the polynomial.

Regarding the problems associated with the camera response recovery, we would recommend to use a linear camera or a modern camera with the RAW output available to capture the HDR images. The usage of RAW data has one additional advantage – the 12-bit precision. The four extra bits in every pixel value result in a lesser quantization error and a wider dynamic range in the input images.

5 Estimate of the Exposure Times

Exposure times reported by the camera on the status display or in the EXIF data are not the times actually used as noted

by Debevec and Malik [3] or Mitsunaga and Nayar [8]. The typical camera provides the exposure times usually rounded to one half stop or one third stop to conform with photographic tradition in marking exposure times. For example, if the camera claims 1/30s or 1/60s was used then the real value could be 1/32s or 1/64s. Even worse, some cheap cameras may have a significant variance in reproducing the same exposure setting. The knowledge of the real exposure times positively affects the radiometric quality of the composed HDR image.

While the produced HDR image is normalized – to have brightness values ranging from 0 to 1 – there is no need to know the absolute value of the exposure times. It is sufficient to estimate the relative exposures against the darkest image. We set the exposure time of the darkest image Δt_1 to be unit, $\Delta t_1 = 1$. The relative exposures of the other images can be computed from the exposure ratios k_j of the consecutive image pairs using formula $\Delta t_j = k_{j-1} \Delta t_{j-1}$.

Recovering exposure times from images themselves can also compensate for uniform changes in illumination between the individual shots provided the change is uniform across the whole scene. The difference in the scene illumination is contained in the recovered exposure times.

5.1 Exposure Ratio Estimate for Linear Camera

For images captured with a nonlinear response camera, Mitsunaga and Nayar [8] described an iterative scheme recovering both – the camera response function and the exposure ratios.

When capturing with a linear response camera there is no need for response recovery. Hence, we use a different approach to estimate the exposure ratio from the captured images. Unlike the method of Mitsunaga and Nayar our method does not require the knowledge of rough exposure estimates. It only needs the brightness transfer function which can be recovered using the image histograms only [4].

In our approach, we fit the brightness transfer function using a linear function. The derivative of this linear function is the unknown exposure ratio. If the camera has a linear response then the brightness values of the image pair are related by the equations

$$\begin{aligned} k_j M_A &= T_j(M_A) \\ M_B &= k_j T_j^{-1}(M_B), \end{aligned} \quad (3)$$

where k_j is the exposure ratio of j -th image pair. T_j is the brightness transfer function relating brightness values M_A of the image j to the brightness values of the image $j + 1$. T_j^{-1} is the inverse brightness transfer function relating brightness values M_B of the image $j + 1$ to the brightness values of the image j .

From (3), $M_A \in \{M: T_j(M) < Z\}$ and $M_B \in \{M: M < Z\}$, where Z is some saturation level, we can construct system S_0 of equations linear in k_j

$$\begin{aligned} k_j M_A \sqrt{h_j(M_A)} &= T_j(M_A) \sqrt{h_j(M_A)} \\ M_B \sqrt{h_{j+1}(M_B)} &= k_j T_j^{-1}(M_B) \sqrt{h_{j+1}(M_B)}. \end{aligned} \quad (4)$$

We used square roots of image histograms h_j, h_{j+1} to weigh the solution of the equations system in the least square sense.




Figure 4: Plots show a linear fit to the blue painted brightness transfer function (BTF). The red line shows initial estimate, the green line (it is covered by the black line in the left plot) shows the result after a final iteration with outliers removed. The left plot is BTF of the two darkest synthetic images. The real BTF used to generate these images is shown in black. The right plot is a real world example. The black line expresses BTF as it would correspond to exposure times reported by the camera.

We solved the linear system in Matlab. Because of noise in the data, there are outliers affecting the least squares fit. An iterative scheme was applied to remove the outliers.

First, the initial ratio $k_j^{(0)}$ is estimated using the system S_0 . Next, the following steps are repeated until the change of k_j is smaller than the required precision. (1) The average square error V_i of the solution S_i is computed. (2) Using current estimate $k_j^{(i)}$, the square errors of all equations in the system S_0 are computed. (3) Equations having errors smaller than $3V_i$ – the inliers – are incorporated in the next linear system S_{i+1} . (4) The system S_{i+1} is used to estimate the next exposure ratio $k_j^{(i+1)}$.

5.2 Limitations of Exposure Times Recovery

The process of exposure times recovery requires that the images sufficiently overlap in the brightness domain. If there is no information to relate brightness values in one image to the brightness values in the other image then the exposure ratio of the two images cannot be estimated.

5.3 Experimental Verification

The test using both the synthetic and the real images was conducted to verify our algorithm estimating exposure ratio of the two differently exposed images captured with a linear response camera.

For the synthetic test, we created six virtual grayscale photographs from an HDR image using the similar approach to the one described in Section 4.2. We used a linear response function instead of a third order polynomial. Similarly we quantized brightness values to 256 discrete levels and added normally distributed noise with mean zero and unit variance.

The left plot in Figure 4 illustrates the estimate of the exposure ratio for the two darkest synthetic images. The darkest images were chosen because they suffer from noise more than brighter images. The initial guess was 3.1378. After seven iterations removing outliers the estimate converged to the value 4.0178. The real exposure ratio used to generate the images was 4.

The right plot in Figure 4 illustrates the estimate of the exposure ratio of two real, differently exposed photographs. The initial guess was 2.1406. It converged to the value 2.1868 after four iterations. The ratio computed from the

exposure times reported by the camera was 1.8750. It can be seen that the estimated value fits the brightness transfer function much better than the value computed from camera data.

6 Image Registration

6.1 The Algorithm

The image registration algorithm is necessary in the case of image capturing with a hand held camera is considered.

Before the registration itself, the camera response function and exposure times are estimated from the histograms of unregistered images.

Only consecutive image pairs are registered. To warp disjoint images a composed transformation is created from the transformations coupling the registered pairs. At the start, each image pair is linearized using the estimated camera response function. Next, the darker image is multiplied by the exposure ratio of the two images. Values above some saturation level are discarded in both images. Similarly values below some noise level are discarded in both images.

The rough estimate of the image shift $(\Delta x_0, \Delta y_0)$ at a pixel precision is found using correlation in Fourier space

$$(\Delta x_0, \Delta y_0) = \max(\mathcal{F}^{-1}(\mathcal{F}(I_1) \cdot \overline{\mathcal{F}(I_2)})), \quad (5)$$

where I_1 and I_2 are the two images in the pair and \mathcal{F} is the complex Fourier spectrum.

The estimates $(\Delta x_0, \Delta y_0)$ and angle $\varphi_0 = 0$ are used to initialize the local search for the image shift $(\Delta x, \Delta y)$ and rotation φ with subpixel precision. The search is performed by the Matlab function `fminsearch` which uses the iterative simplex search method. The optimized criterion ε is the square of normalized difference summed across all corresponding pixel candidates

$$\varepsilon_i = \sum_{x,y} \left(\frac{I_1(x,y) - W_2(x,y)}{n(x,y)} \right)^2, \quad (6)$$

where W_2 is image I_2 warped using current $(\Delta x_i, \Delta y_i)$ and angle $\varphi_i = 0$ in the iteration i . $n(x,y)$ is the normalizing function balancing the importance of bright and dark pixels. We used the average value of the potentially corresponding pixels as the normalization term

$$n(x,y) = \begin{cases} \frac{I_1(x,y) + W_2(x,y)}{2} & \text{for } \frac{I_1(x,y) + W_2(x,y)}{2} \geq t \\ t & \text{for } \frac{I_1(x,y) + W_2(x,y)}{2} < t \end{cases}, \quad (7)$$

where t is normalization threshold preventing the boost of importance of the very dark noisy pixels.

6.2 Experimental Verification

Figure 5 demonstrates the performance of the implemented registration algorithm. The left picture shows a tonemapped version of the HDR image captured with a hand held camera. The image was stitched from three photographs differing by two stops in exposure. The photographs were captured at a wide end of the zoom lens – a 28mm focal length in 35mm camera equivalent. The use of a short focal length makes the registration process difficult due to the wide field of view and the lens distortion.



Figure 5: The HDR image taken with a hand held camera. (top) mid-exposed photograph from the source image set. (middle) HDR image from photographs registered for shift. (bottom) HDR image from photographs registered for shift and rotation.

The right side of Figure 5 shows enlarged details. Details in the first row were taken from the second source photograph and provide a sharpness comparison with the HDR stitches. Details in the middle row were taken from the HDR image stitched from photographs registered by modeling shift only. The registration process was performed using the correlation in Fourier domain at a pixel precision. Details in the bottom row were taken from the HDR image stitched from photographs registered by shift and rotation modeling at a subpixel resolution. This was done by a non-linear optimization using a simplex method. Details in the first column were taken from the middle of the scene while details in the second column were from the lower left corner.

From the details shown, it can be seen that the registration by modeling shift at a pixel precision is not sufficient. On the other side, the images registered at a subpixel precision by modeling shift and rotation provide a HDR image which is only slightly softer than the source images. The stitched HDR image provides almost the same amount of detail compared to the information contained in the all source images. However, the total amount of detail in the HDR image is much higher when comparing to one source image only. This is caused by the missing information in the overexposed or underexposed areas of the source image.

7 Conclusions

We showed that the recovery of camera response function using the differently exposed images is a difficult task. We implemented state-of-the-art algorithms according to descriptions given by their authors. We learned that they are very sensitive to the noise in data. In agreement with our results, we see the way around the problem in the use of linear 12-bit RAW images which are now offered by many modern mid-class and professional cameras.

We proposed a robust method recovering exposure times from unregistered photographs captured by a linear response camera. The method can be used with linear RAW images. Next, we showed that capturing HDR images is possible

without mounting the camera on a tripod and described a method to register captured images differing in exposure. The presented method expands the usability of HDR widening to hand held photography. Methods were implemented in Matlab and tested on both synthetic and real data.

Acknowledgement

The authors wish to thank Jan Čech for advice with geometric image registration. The authors also acknowledge the critical and constructive comments by the reviewers. The second author was supported by the projects INTAS 04-77-7347, CAK 1M0567 and BeNoGo IST-2001-39184.

References

- [1] A. Adams. *The Negative*. Bulfinch, 1995.
- [2] Young-Chang Chang and J. F. Reid. RGB calibration for color image analysis in machine vision. *IEEE Transactions on image processing*, 5(10):1414–1422, 1996.
- [3] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH*, pages 369–378, 1997.
- [4] M. D. Grossberg and S. K. Nayar. What can be known about the radiometric response function from images? In *European Conference on Computer Vision, Part IV*, pages 189–205, 2002.
- [5] S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High dynamic range video. *ACM Trans. Graph.*, 22(3):319–325, 2003.
- [6] S. Mann, C. Manders, and J. Fung. Painting with looks: Photographic images from video using quantimetric processing. In *ACM international conference on Multimedia*, pages 117–126, 2002.
- [7] S. Mann and R. Picard. Being ‘undigital’ with digital cameras: Extending dynamic range by combining differently exposed pictures. In *IS&T 46th annual conference*, pages 422–428, 1995.
- [8] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Computer Vision and Pattern Recognition*, volume 2, pages 374–380, 1997.
- [9] Y. Y. Schechner and S. K. Nayar. Generalized mosaicing: High dynamic range in a wide field of view. *Computer Vision*, 53(3), 2003.
- [10] IMS Chips HDRC. <http://www.ims-chips.com/home.php3?id=d0841>.
- [11] Kodak KAC-9628. <http://www.kodak.com/global/en/digital/ccd/products/cmos/KAC-9628/indexKAC-9628.jhtml?id=0.1.10.4.11/&lc=en>.
- [12] Pixim. http://www.pixim.com/html/tech_about.htm.
- [13] Silicon Vision. <http://www.siliconvision.de>.
- [14] SmaL Camera. <http://www.smalcamera.com/technology.html>.
- [15] Y. Tsin, V. Ramesh, and T. Kanade. Statistical calibration of CCD imaging process. In *International Conference on Computer Vision*, volume 1, July 2001.
- [16] G. Ward. Fast, robust image registration for compositing high dynamic range photographs from hand-held exposures. *Graphics Tools*, 8(2):17–30, 2003.