High Dynamic Range from Multiple Images: Which Exposures to Combine?*

Michael D. Grossberg and Shree K. Nayar
Department of Computer Science, Columbia University
New York, New York 10027
E-mail: {mdog, nayar}@cs.columbia.edu

Abstract

Many computer vision algorithms rely on precise estimates of scene radiances obtained from an image. A simple way to acquire a larger dynamic range of scene radiances is by combining several exposures of the scene. The number of exposures and their values have a dramatic impact on the quality of the combined image. At this point, there exists no principled method to determine these values. Given a camera with known response function and dynamic range, we wish to find the exposures that would result in a set of images that when combined would emulate an effective camera with a desired dynamic range and a desired response function.

We first prove that simple summation combines all the information in the individual exposures without loss. We select the exposures by minimizing an objective function that is based on the derivative of the response function. Using our algorithm, we demonstrate the emulation of cameras with a variety of response functions, ranging from linear to logarithmic. We verify our method on several real scenes. Our method makes it possible to construct a table of optimal exposure values. This table can be easily incorporated into a digital camera so that a photographer can emulate a wide variety of high dynamic range cameras by selecting from a menu.

1 Capturing a Flexible Dynamic Range

Many computer vision algorithms require accurate estimates of scene radiance such as color constancy [9], inverse rendering [13, 1] and shape recovery [17, 8, 18]. It is difficult to capture both the wide range of radiance values real scenes produce and the subtle variations within them using a low cost digital camera. This is because any camera must assign a limited number of brightness values to the entire range of scene radiances. The response function of the camera determines the assignment of brightness to radiance. The response therefore determines both the camera's sensitivity to changes in scene radiance and its dynamic range.

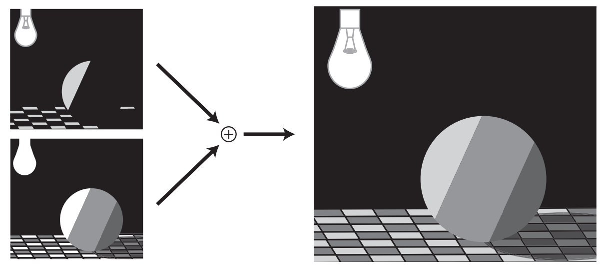

(a) Small and large exposures combine to capture a high dynamic range

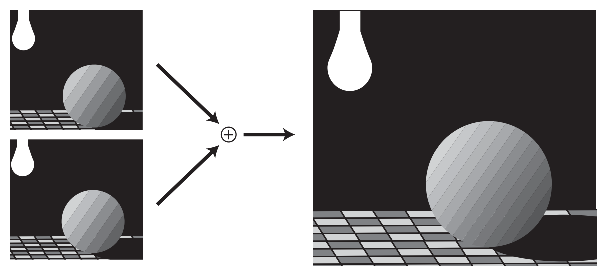

(b) Similar exposures combine to capture suble variations

Figure 1: Illustration showing the impact of the choice of exposure values on which scene radiances are captured. (a) When large and small exposures are combined the resulting image has a high dynamic range, but does not capture some scene variations. (b) When similar exposure values are combined, the result includes subtle variations, but within a limited dynamic range. In both cases, a set of exposures taken with a camera results in an "effective camera." **Which exposures must we use to emulate a desired effective camera?**

A simple method for extending the dynamic range of a camera is to combine multiple images of a scene taken with different exposures [6, 2, 3, 10, 11, 12, 15, 16]. For example, the left of Fig. 1(a) shows a small and a large exposure, each capturing a different range of scene radiances. The illustration on the right of Fig. 1(a) shows that the result of combining the exposures includes the entire dynamic range of the scene.

*This work was completed with support from a National Science Foundation ITR Award (IIS-00-85864) and a grant from the Human ID Program: Flexible Imaging Over a Wide Range of Distances Award No. N000-14-00-1-0929

Note that by using these exposures values we fail to capture subtle variations in the scene, such as the shading of the ball. Once these variations are lost they can not be restored by methods that change the brightness of an image, such as the recent work on tone mapping [4, 5, 14].

In Fig. 1(b), two similar exposures combine to produce an image that captures subtle variations, but within a limited dynamic range. As a result, in both Fig. 1(a) and (b), the images on the right can be considered as the outputs of two different "effective cameras." The number and choice of exposures determines the dynamic range and the response of each effective camera. This relationship has been ignored in the past. In this paper we explore this relationship to address the general problem of determining which exposure values to use in order to emulate an effective camera with a desired response and a desired dynamic range. Solving this problem requires us to answer the following questions:

- **How can we create a combined image that preserves the information from all the exposures?** Previous work suggested heuristics for combining the exposures [3, 11, 12]. We *prove* that even without linearizing the camera, simple summation preserves all the information contained in the set of individual exposures.

- **What are the best exposure values to achieve a desired effective response function for the combined image?** It is customary to arbitrarily choose the number of exposures and the ratio (say, 2) between consecutive exposure values [3, 10, 11, 12]. For example, when this is done with a linear real camera, the resulting combined image is relatively insensitive to changes in large radiances. This can bias vision algorithms that use derivatives of radiance. Such biases are eliminated using our algorithm, which selects the exposure values to best achieve a desired response.

- **How can we best achieve a desired dynamic range and effective response function from a limited number of images?**

  It is common to combine images with consecutive exposure ratios of 2 (see [3, 11, 12]). to create a high dynamic range image. With that choice of exposure ratio, is often necessary to use 5 or more exposures to capture the full dynamic range of a scene. This is impractical when the number of exposures that can be captured is limited by the time to acquire the images, changes in the scene, or resources needed to process the images. Our algorithm determines the exposure values needed to best emulate a desired camera with a fixed number of images.

Our method allows us to emulate cameras with a wide variety of response functions. For the class of linear real cameras, we present a table of optimal exposure values for emulating high dynamic range cameras with, for example, linear and logarithmic (constant contrast) responses. Such a table can be easily incorporated into a digital camera so that a photographer can select his/her desired dynamic range and camera response from a menu. In other words, a camera with fixed response and dynamic range can be turned into one that has a "flexible" dynamic range. We show several experimental results using images of real scenes that demonstrate the power of this notion of flexible dynamic range.

2 The Effective Camera

When we take multiple exposures of the same scene, each exposure adds new information about the radiance values in the scene. In this section, we create an effective camera by constructing a single image which retains all the information from the individual exposures.[1] By information we mean image brightness values which represent measurements of scene radiance.

Scene radiance is proportional to image irradiance E [7]. In a digital camera, the camera response function f jumps from one image brightness value B to the next at a list of positive irradiance values (shown below the graph in Fig. 2) which we call the *measured irradiance levels*. An image brightness value indicates that the corresponding measured irradiance lies in the interval between two of these levels. Hence, without loss of generality, we define B as the index of the first of these two levels, E_B, so that $f(E_B) = B$. Hence, the response function is equivalent to the list of measured irradiance levels.[2]

Now, consider the measured irradiance levels using unit exposure $e_1 = 1$ with a real non-linear camera having 4 brightness levels.[3] These levels are shown on the bar at the bottom of Fig. 3(a). The irradiance levels for a second exposure scale by $1/e_2$, as shown in Fig. 3(b). We combine the measured irradiance levels from the first and the second exposures by taking the *union of all the*

[1] The value we call *exposure* accounts for all the attenuations of light by the optics. One can change the exposure by changing a filter on the lens, the aperture size, the integration time, or the gain.

[2] Note that the slope of the response function determines the density of the levels, as shown by the short line segment in Fig. 2.

[3] Note that the number of exposures and brightness levels are for illustration only. Our arguments hold in general.

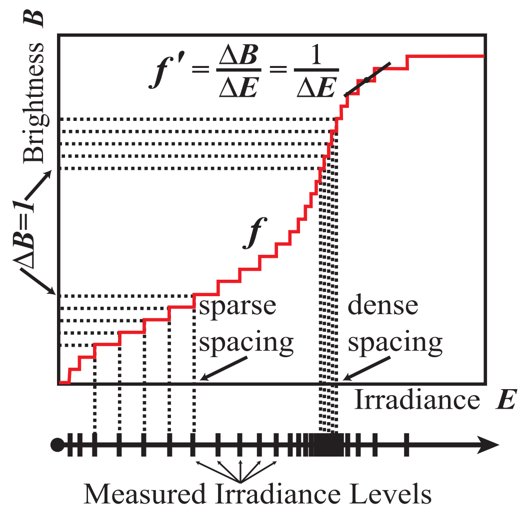

Figure 2: The camera response function f relates irradiance E to image brightness B. In a digital camera, the discontinuities of this function define a list of measured irradiance levels shown on the bar beneath the graph. The slope of the camera response function (short line segment) gives the density of these levels.

levels, shown under the graph in Fig. 3(c). This combined list of levels uniquely determines the response h shown in Fig. 3(c). It is important to note that the response h may also be obtained by summing the response functions[4] $f(E)$ in Fig. 3(a) and $f(e_2\ E)$ in Fig. 3(b). Therefore, for n exposures, the combined camera response is

$$h(E) = \sum_{j=1}^{n} f(e_j\ E). \tag{1}$$

In terms of images, let I_E be a vector of image irradiance values. Then the jth captured image of brightness values is $f(e_j\ I_E)$ where the real non-linear response is applied to each pixel. Hence, Eq. 1 implies that summing the images results in an image $h(I_E)$. Since the response h is the same as obtained by taking the union of the measured irradiance levels from all the exposures, we have therefore proven the following:[5]

Theorem 1 *The sum of a set of images of a scene taken at different exposures includes all the information in the individual exposures.*

[4]Here we assume that none of the irradiance levels in different exposures*exactly* coincide. In practice this is not a problem.

[5]We emphasize that we do not assume the camera is linear nor do we assume that the camera is linearized. We assume only that the non-linear response f is monotonic.

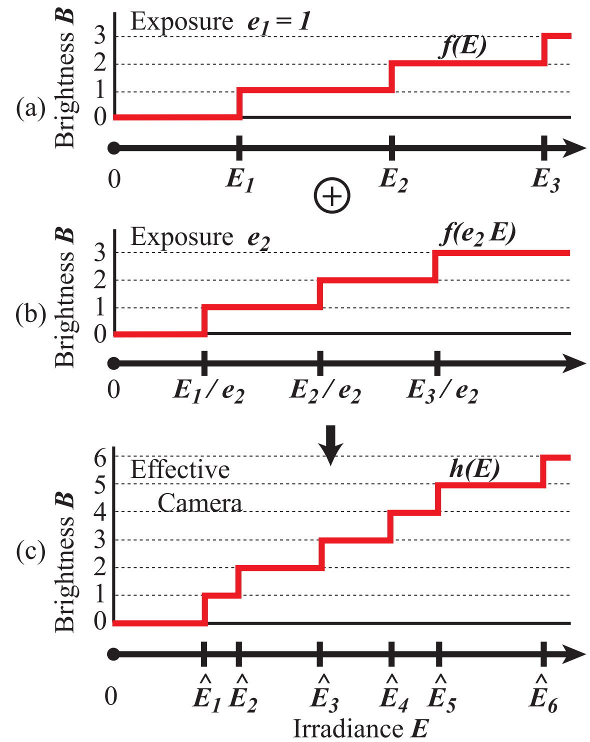

Figure 3: (a) The response function using unit exposure is equivalent to a list of irradiance levels where the brightness changes. These levels are shown on the bar beneath the graph. (b) In a second exposure, the levels in (a) scale by $(1/e_2)$. (c) The irradiance levels for the effective camera are obtained from the union of the levels in (a) and (b) which determines the response function h. Note that this response can also be obtained by summing the individual response functions in (a) and (b). As a result, simple summation of the acquired images results in an effective camera image that preserves all the information in the individual images.

3 Finding the Best Exposures

We now determine the best exposures for an effective camera, given the response of the real camera and a desired response function.[6] To do this, we first propose an objective function which measures how well an effective camera emulates a desired camera. We then find the exposures that minimize the objective function.

3.1 The Objective Function

A naïve measurement of similarity between a desired response and the response achieved by an effective camera is the norm of their difference. This difference, however, is a poor measure of similarity. For example, adding a constant brightness to one response function will change this measure. Nevertheless, it does not change the func-

[6]The desired response determines the desired dynamic range.

tion's discontinuities and thus does not change its measured irradiance levels. The distances between the irradiance levels determine the accuracy in measuring the irradiances. We determine these distances using the following observation (illustrated in Fig. 2) which forms the basis of our objective function.

Observation 1 *The derivative of the camera response determines the inverse of the distances between irradiance levels.*

Recall that the response h of an effective camera depends on the response f of the real camera, the number n of exposures, and a vector of exposure values $\mathbf{e} = (e_1, e_2, \ldots, e_n)$ (see Eq. 1). Let h' be the derivative of the effective response, let g' be the derivative of the desired response, and let p be a positive number indicating the norm we use. Typically, p is 1 or 2. We define the objective function as:

$$\xi(n, \mathbf{e}) = \int_{E_{\min}}^{E_{\max}} |g' - h'|^p \ w \ dE, \tag{2}$$

where $E_{\min}$, $E_{\max}$ is the interval where g' is non-zero and w is a weighting function defined as

$$w(E) = \begin{cases} 0 & \text{when } g'(E) < h'(E) \\ 1 & \text{otherwise.} \end{cases} \tag{3}$$

This weighting function prevents penalizing the response h of the effective camera for levels spaced more *densely* than required[7] by the desired response g.

We change the objective function ξ in two ways to account for camera noise. We modify the response of the effective camera by removing levels most affected by noise and we add a term $V(n, \mathbf{e})$ to ξ that penalizes exposures which increase the noise variance. As an example, if we model irradiance E with Gaussian noise having variance σ^2 added and assume a linear camera, then $V(n, \mathbf{e}) = \sigma^2 \sum_{i=1}^{n} 1/e_i$.[8] With the inclusion of this noise term, the objective function becomes:

$$\tilde{\xi}(n, \mathbf{e}) = \xi(n, \mathbf{e}) + cV(n, \mathbf{e}), \tag{4}$$

where the constant c weights the noise term. This objective function provides a simple means to evaluate how well the effective camera resulting from a set of exposures emulates the desired camera.

3.2 Optimal Exposures

For a fixed number of exposures, we find the exposure values that minimize the objective function of Eq. 4. We constrain the exposure values to be positive and assume them to be in increasing order. It is easy to impose these constraints. One difficulty, however, is that the objective function is not continuous since w in Eq. 3 is not continuous. This discontinuity makes the objective function difficult to minimize. We have yet to find the best means to minimize the objective function. For now, we have implemented a simple exhaustive search. To find the minimum number of exposures, we iterate the search starting with 2 exposures and increasing the number of exposures by one in each iteration. We continue until the minimum of the objective function falls below a given tolerance.

The search efficiency is *not* critical since we can generate a table for a large number of desired effective cameras off-line. As an example, the table shown in Figure 4(a) was generated using our method assuming a real 8-bit linear camera. This table gives the exposures for emulation of a linear, a gamma = 1/2, and a logarithmic (constant contrast) response. Such a table can be integrated into a digital camera. The interface of the camera would allow a user to choose a desired response and desired dynamic range from a menu, as shown in Figure 4(b). The camera would automatically capture the appropriate sequence of images by obtaining the exposure values from the table.

We see from the table in Figure 4(a) that the algorithm creates a linear effective response by selecting exposures which are similar. As levels from different exposures interleave, the density of the levels increases within a limited dynamic range. The dynamic range of the constant contrast case is extended with successively larger exposures, as seen in the table. To compute these exposures we ignored noise, set p to 1 in Eq. 2, and normalized the exposures so that the first exposure is 1.

Some of the computed exposures are close together, requiring 4 significant digits to show the differences. When this is not possible, a rough approximation to the best exposures will often provide significant benefit. Even with a low cost consumer camera it is possible to achieve a large number of possible exposures by varying the settings of integration time, aperture, and gain.[9] Note that we can search directly over the pos-

[7] The density of the levels can always be decreased by dropping levels. Dropping levels from the response of the effective camera decreases its derivative h', thus changing its shape. For example, the response of an effective camera constructed from a *linear* real camera is always convex down. By decreasing h', we can make the response convex up.

[8] The total variance of the noise depends on the kind of exposure change (gain, integration time, or aperture) as well as the response of the camera.

[9] The illumination is assumed to be constant. Small uniform

Effective Camera	Number of Exposures	Exposures 1	2	3	4	5	6
Linear	2	1	1.003				
	3	1	1.003	2.985			
	4	1	1.015	1.019	3.672		
	5	1	1.003	1.007	1.166	4.975	
	6	1	1.003	1.007	1.031	1.078	5.636
Gamma = 1/2	2	1	3.094				
	3	1	1.003	5.146			
	4	1	1.003	3.019	11.23		
	5	1	1.003	1.006	5.049	18.56	
	6	1	1.003	1.006	2.866	8.564	33.7
Constant Contrast (log)	2	1	20.24				
	3	1	9.91	88.38			
	4	1	4.689	37.23	280		
	5	1	4.831	29.18	144.9	763.5	
	6	1	3.979	16.01	64.22	305.4	1130

(a)

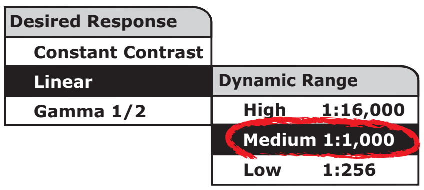

(b)

Figure 4: (a) A table showing exposure values for emulating a few effective cameras. The exposures have been normalized so that the smallest exposure is 1. (b) A menu interface which could be included in a digital camera. The digital camera would lookup the sequence of exposure values corresponding to the selection, capture a sequence of images with these exposures, combine them, and output the resulting high dynamic range image.

sible settings available for a particular camera to best achieve a desired effective camera for a fixed number of exposures.[10]

4 Experiments

We created four effective cameras with our method, using both linear and non-linear real cameras. Like any real camera, we can adjust the exposure of our effective camera. We do this by varying the shortest exposure. For each effective camera we combined 3 exposures. As a baseline, we evaluated our results against those obtained with the commonly used exposures 1, 2, and 4, [3, 11, 12]. We show our method decreases or eliminates the posterization[11] that occurs when brightness levels are too widely spaced to capture differences in the scene radiance.

changes in illumination change the effective exposure. Since the response of the real camera is known such changes may be estimated from the images.

[10] Instead of fixing the number of exposures, we may fix the total time required to take all the exposures.

[11] Many image processing programs, such as Adobe Photoshop, have a "posterize" function. By reducing the number of brightness levels, an image takes on the appearance of an art poster.

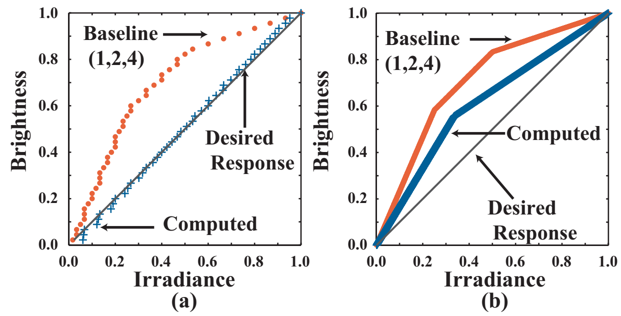

Figure 5: Emulation of a linear camera with high dynamic range. (a) A graph showing the response functions for two effective cameras obtained from a 4-bit linear real camera. The exposures computed by our algorithm (1, 1.05, 1.11) produce an effective camera closer to the desired linear response, than produced with the 3 baseline exposures (1, 2, 4). (b) The effective response obtained from an 8-bit real camera by our algorithm is more complex than in (a) but it results in smaller errors than obtained with the baseline exposures.

4.1 Adding Bits to a Linear Camera

The dynamic range of a linear real camera can be extended by combining exposures whose values are powers of 2. This results in a non-linear effective response, similar to a gamma curve [10]. Yet, a *linear* effective response is often desirable for vision applications because it provides a uniform assignment of image brightness to scene radiance. We start with the simple example of combining 3 exposures from a linear real camera with 4 bits. Our algorithm specifies exposures close together (1, 1.05, 1.11). The response from the resulting effective camera is shown with the response resulting from the baseline exposures (1, 2, 4) in Fig. 5(a). Using the exposures computed with our algorithm allows us to achieve a more uniform response. This is especially true for the higher irradiance values as compared to what is achieved with the baseline exposures.

For an 8-bit linear camera, the algorithm determines two exposures close together, 1 and 1.003, and a third exposure which is larger than the other two, 2.985. Fig. 5(b) shows the improved uniformity of the irradiance levels. Our method decreases the error when compared with using the baseline exposures.

Fig. 6(a) shows the wide spacing of irradiance levels in an image of a synthetic linear ramp obtained from a 4-bit camera by combining images with exposures (1,2,4). Fig. 6(b) shows that our method reduces the spacing of irradiance levels. Fig. 6(c) shows the same exposure values used in Fig. 6(a) but with Gaussian noise with a standard deviation of 1.95 brightness levels added to the input images. The wide spacing of the irradiance levels remains. In Fig. 6(d), the narrower spacing of levels from Fig. 6(b) results in a smoother gradient with the same noise as used in Fig. 6(c). Fig. 6(e) shows the result obtained from a 12-bit camera image of the linear ramp for comparison.

The baseline exposures (1,2,4) provide relatively few irradiance levels for large irradiance values. This is clearly demonstrated by our experimental result shown in Fig. 7. The acquired images come from a real 4-bit/channel camera simulated by dropping bits from a 12-bit/channel camera. The image in Fig. 7(a) was obtained by summing three images taken with the baseline exposures. Due to the wide spacing of irradiance levels the resulting image has posterization artifacts. The image in Fig. 7(b) was obtained directly from a 12-bit/channel linear camera. Fig. 7(c) shows the image obtained by combining exposures chosen using our algorithm. We see that our image is quite close to the ground truth.

4.2 Constant Contrast Effective Camera

The human visual system is more sensitive to contrast than to absolute differences in irradiance. For differences in contrast to be independent of absolute irradiance, the measured irradiance levels must have constant contrast spacing. This amounts to a desired response that has the form of a log curve $g(E) = \alpha \log_2(E) + \beta$. The response function of the effective camera created from the baseline exposures (1, 2, 4) does not provide enough irradiance levels at low irradiances to achieve constant contrast spacing, as is apparent from Fig. 8 (a). The exposures determined by our algorithm (1, 9.91, 88.38) give an effective camera whose response is much closer to the desired response. In Fig. 8 (b) we see that as we increase the number of exposures from 3 to 5, our response rapidly approaches the desired response.

Fig. 9 shows a tile floor with a strong illumination gradient. We dropped bits from a linear camera (12-bit) to make the comparison more apparent in the images, simulating a 6-bit camera. The image in Fig. 9(a) was obtained from adding the baseline exposures. It does not properly capture the variations in the scene radiance. Summing the exposures given by our algorithm results in the image shown in Fig. 9(b) which does not have the noise and posterization in Fig. 9(a).

4.3 Using Non-linear Real Cameras

Many consumer 8-bit cameras extend their dynamic range by unevenly spacing the irradiance levels with a non-linear response. For vision applications requiring a

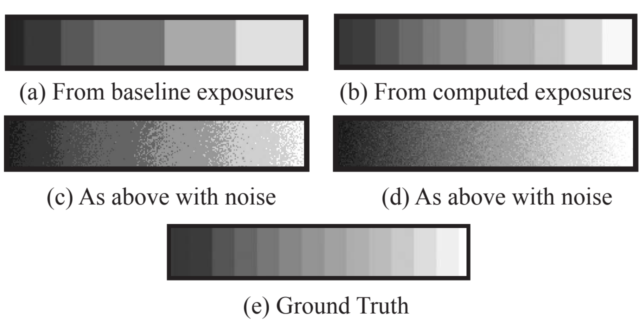

Figure 6: Validation using synthetic images. (a) An image of a linear synthetic ramp obtained by combining 4-bit images with baseline exposures showing wide spacing of irradiance levels. (b) An image obtained from combining images with exposures determined by our algorithm (1, 1.05, 1.11) with narrower spacing of irradiance levels. (c) When noise is added to the input images used in (a), the wide spacing is still apparent. (d) When noise is added to the same exposures as in (b) the ramp appears almost continuous. (e) An image from a 12-bit linear camera for comparison.

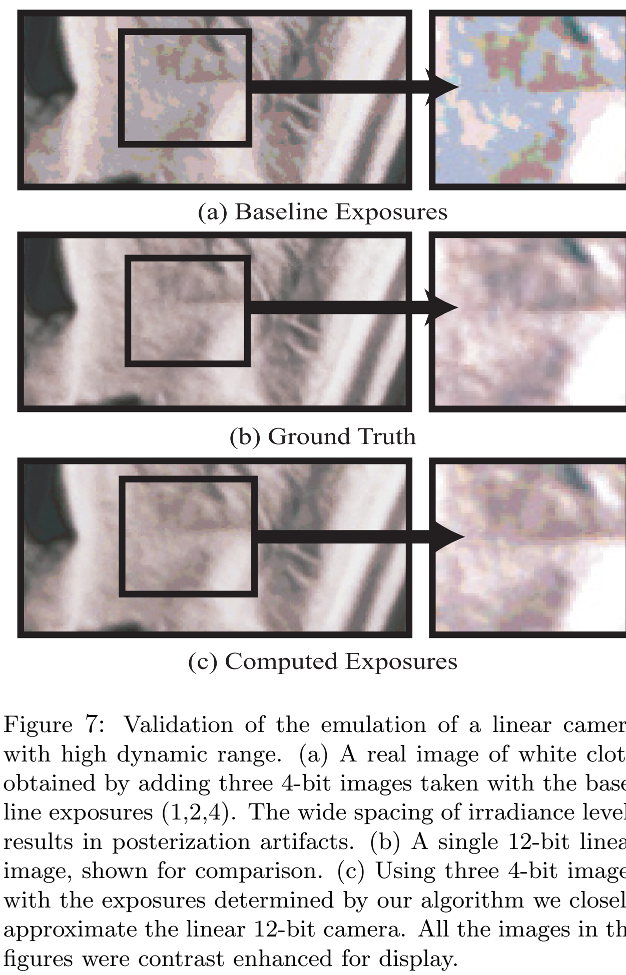

Figure 7: Validation of the emulation of a linear camera with high dynamic range. (a) A real image of white cloth obtained by adding three 4-bit images taken with the baseline exposures (1,2,4). The wide spacing of irradiance levels results in posterization artifacts. (b) A single 12-bit linear image, shown for comparison. (c) Using three 4-bit images with the exposures determined by our algorithm we closely approximate the linear 12-bit camera. All the images in the figures were contrast enhanced for display.

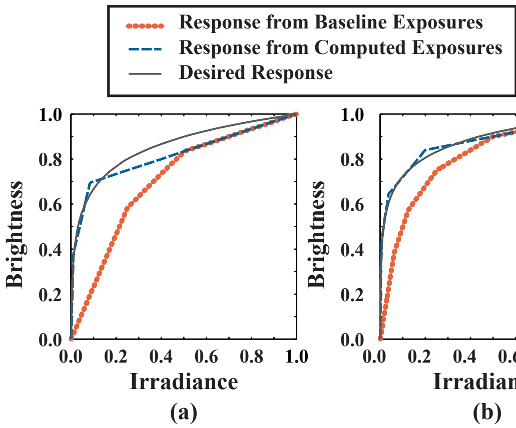

Figure 8: Emulation of a constant contrast camera. (a) The response obtained by a linear real camera using the baseline exposures clearly has too few levels for low irradiance values as compared with the constant-contrast response. Our algorithm determines three exposures (1,9.91,88.38) so that the response of the effective camera matches the desired curve more closely. (b) The same as (a) but with 5 exposures used.

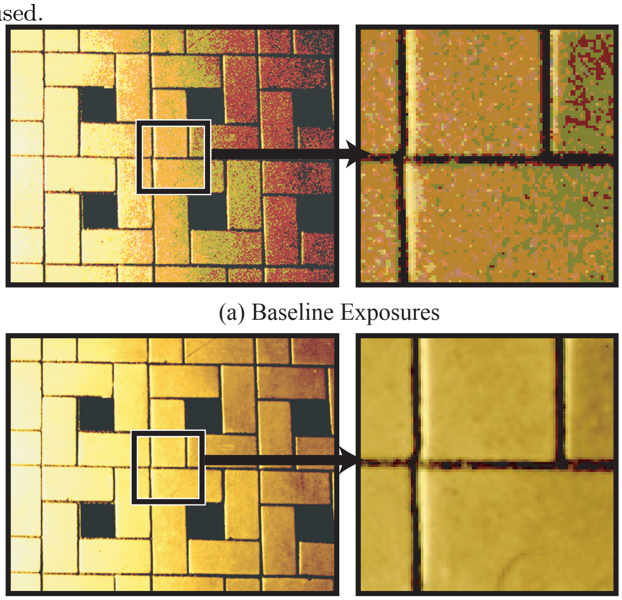

(a) Baseline Exposures

(b) Computed Exposures

Figure 9: Validation of the emulation of a constant contrast camera. (a) Image resulting from adding three 6-bit images taken with the baseline exposures. The posterization due to the wide spacing of irradiance levels is clearly visible. (b) Images taken using exposures determined by our algorithm provide a much denser spacing of irradiance levels eliminating the posterization.

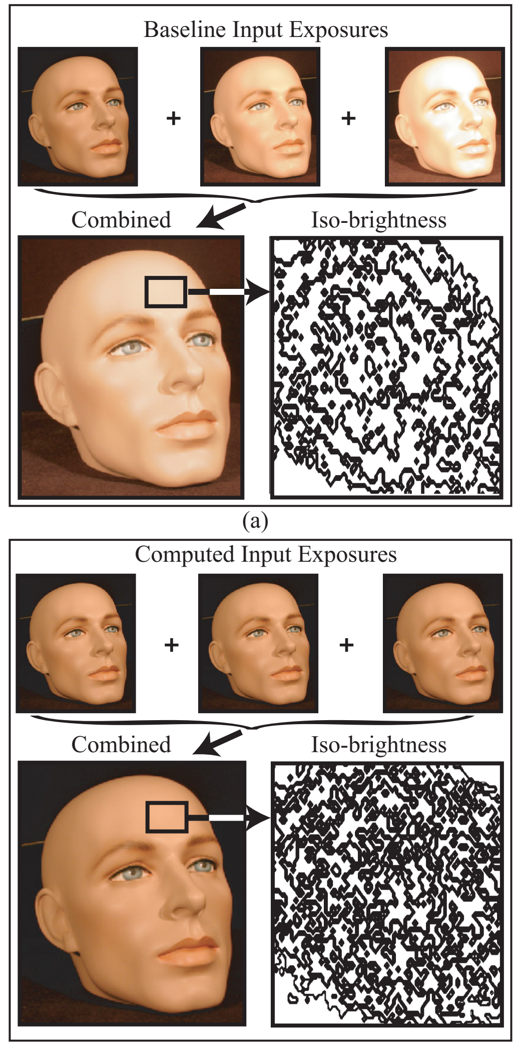

Figure 10: (a) Three 8-bit images of a limited dynamic range scene using the baseline exposures (1, 2, 4) are combined. The resulting image has widely spaced irradiance levels in the highlights as shown by the iso-brightness contours of the detail. (b) By combining three similar exposures (1, 1.0073, 1.0493) chosen by our algorithm we decrease the distance between irradiance levels as seen in the denser iso-brightness contours.

linear response, simply changing image brightness values (e.g. linearizing) cannot help when the spacing of irradiance levels is too wide to capture scene radiance variations, as shown in Fig. 1(a). By taking several images with exposures chosen by our algorithm we can narrow the spacing of irradiance levels to better approximate the uniform irradiance levels of a linear response, as illustrated in Fig. 1(b).

We verified this application with an 8-bit Nikon 990 Coolpix camera, whose response we calibrated with a Macbeth chart. The widely spaced irradiance levels from the baseline exposures result in widely spaced iso-brightness curves around the highlights in the image shown in Fig. 10(a). By combining our computed exposures we obtain greater sensitivity resulting in much more closely spaced iso-brightness curves Fig. 10(b).

5 Conclusion

We proved that when multiple exposures are combined by simple summation we preserve the information from each of the exposures. We also provided a simple formula for the response function of the effective camera that results from this summation. We presented an objective function for the error between the response function of the effective camera and a desired response function based on their derivatives. We used a simple search to obtain exposures that, when combined, best achieve a desired response function. This search is performed off-line to produce a list of exposures to combine.

Using our method, we generated a table which lists the exposures needed to emulate cameras with linear, gamma = 1/2, and log (constant contrast) response functions. Hence, is possible to emulate a wide selection of commonly used response functions. Moreover, the same framework could also be used to emulate image *dependent* response functions, for example based on the histogram equalization of an initial image.

We verified our method using images of real scenes. We have shown that our algorithm provides a means to construct an effective camera with a flexible camera response function and selectable dynamic range. With our method, photographers and scientists can emulate the high dynamic range camera that meets their needs with a low cost consumer camera.

References

[1] R. Basri and D.W. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218–233, February 2003.

[2] P.J. Burt and R.J. Kolczynski. Enhanced Image Capture Through Fusion. In *Proc. of the ICCV*, pages 173–182, 1993.

[3] P. E. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. of ACM SIGGRAPH*, pages 369–378, 1997.

[4] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high dynamic range images. In *Proc. of ACM SIGGRAPH*, pages 369–378, 2002.

[5] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. In *Proc. of ACM SIGGRAPH*, pages 249–256, 2002.

[6] R. Ginosar and Y. Y. Zeevi. Adaptive Sensitivity/intelligent scan image processor. In *SPIE Visual Communications and Image Processing*, volume 1001, November 1988.

[7] B. K. P. Horn. *Robot Vision.* MIT Press, Cambridge, MA, 1986.

[8] B.K.P. Horn and M.J. Brooks. *Shape from Shading.* MIT Press, 1989.

[9] E.H. Land and J.J. McCann. Lightness and Retinex Theory. *JOSA*, 61(1):1–11, 1971.

[10] B.C. Madden. Extended Intensity Range Image. In *Technical Report 366, Grasp Lab,UPenn*, 1993.

[11] S. Mann and R. Picard. Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. In *In Proceedings of IS& T, 46th annual conference*, pages 422–428, 1995.

[12] T. Mitsunaga and S. K. Nayar. Radiometric Self Calibration. In *Proc. of CVPR*, volume 2, pages 374–380, June 1999.

[13] R. Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *PAMI*, 24(10):1322–1333, October 2002.

[14] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. In *Proc. of ACM SIGGRAPH*, pages 369–378, 2002.

[15] M.A. Robertson, S. Borman, and R.L. Stevenson. Estimation Theoretic Approach to Dynamic Range Improvement Through Multiple Exposures. In *Proc. of ICIP*, pages 159–163, 1999.

[16] Y.Y. Schechner and S.K. Nayar. Generalized Mosaicing. In *Proc. of the ICCV*, pages I: 17–24, 2001.

[17] R.J. Woodham. Photometric Method for Determining Surface Orientation from Multiple Images. *Optical Engineering*, 19(1):139–144, January 1980.

[18] T. Zickler, P.N. Belhumeur, and D.J. Kriegman. Helmholtz Stereopsis: Exploiting Reciprocity for Surface Reconstruction. In *Proc. of the ECCV*, page III: 869 ff., 2002.