

# HDR Image Capture

## 04

HDR images may be captured from real scenes or rendered using 3D computer graphics (CG) techniques such as radiosity and raytracing. A few modern graphics cards are even capable of generating HDR images directly. The larger topic of CG rendering is well covered in other textbooks [24,37,58,117,144]. In this chapter, the focus is on practical methods for capturing

high-quality HDR images from real scenes using conventional camera equipment. In addition, commercial hardware designed to capture HDR images directly is beginning to enter the market, which is discussed toward the end of this chapter.

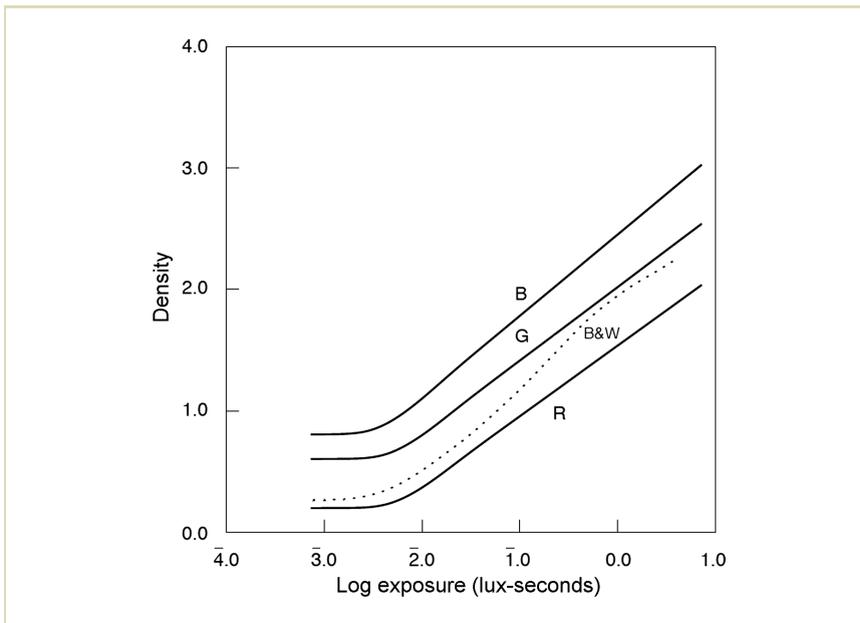
### 4.1 PHOTOGRAPHY AND LIGHT MEASUREMENT

A camera is essentially an imperfect device for measuring the radiance distribution of a scene, in that it cannot capture the full spectral content and dynamic range. (See Chapter 2 for definitions of color and radiance.) The film or image sensor in a conventional or digital camera is exposed to the color and dynamic range of a scene, as the lens is a passive element that merely refocuses the incoming light onto the image plane. All of the information is there, but limitations in sensor design prevent cameras from capturing all of it. Film cameras record a greater dynamic range than their digital counterparts, especially when they expose a negative emulsion.

Standard black-and-white film emulsions have an inverse response to light, as do color negative films. Figure 4.1 shows example response curves for two film emulsions, demonstrating a sensitive range of nearly 4 log units, or a 10,000:1

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35



**FIGURE 4.1** Characteristic film curves showing response for color (R, G, and B) and black-and-white negative films over nearly four orders of magnitude.

contrast ratio. Depending on the quality of the lens and the blackness of the camera's interior, some degree of "flare" may inhibit this range, particularly around a bright point such as the sun or its reflection. Although the capability of recording 4 log units of dynamic range is there, flare may reduce the effective dynamic range somewhat.

The film development process may also limit or enhance the information retrieved from the exposed emulsion, but the final constraining factor is of course the printing process. It is here where tone mapping takes place, in that the effective dynamic range of a black-and-white or color print is about 100:1 at best. Darkroom

1 techniques such as dodge-and-burn may be used to get the most out of a negative, 1  
 2 although in an industrial film processing lab what usually happens is more akin to 2  
 3 autoexposure after the fact. 3

4 To extract the full dynamic range from a negative, we need to digitize the de- 4  
 5 veloped negative or apply a “dry developing method” such as that developed by 5  
 6 Applied Science Fiction and marketed in the Kodak Film Processing Station [25]. 6  
 7 Assuming that a developed negative is available, a film scanner would be required 7  
 8 that records the full log range of the negative in an HDR format. Unfortunately, no 8  
 9 such device exists, although it is technically feasible. However, one can take a stan- 9  
 10 dard film scanner, which records either into a 12-bit linear or an 8-bit sRGB color 10  
 11 space, and use multiple exposures to obtain a medium-dynamic-range result from a 11  
 12 single negative. The process is identical to the idealized case for multiple-exposure 12  
 13 HDR capture, which we describe in the following section. The same method may 13  
 14 be used to obtain an HDR image from a sequence of exposures using a standard 14  
 15 digital camera, or to enhance the dynamic range possible with a film camera. 15  
 16

17  
 18 **4.2 HDR IMAGE CAPTURE FROM MULTIPLE** 18  
 19 **EXPOSURES** 19  
 20

21 Due to the limitations inherent in most digital image sensors, and to a lesser degree 21  
 22 in film emulsions, it is not possible to capture the full dynamic range of an image 22  
 23 in a single exposure. However, by recording multiple exposures a standard camera 23  
 24 with the right software can create a single HDR image (i.e., a *radiance map*, as defined 24  
 25 in Chapter 2). These exposures are usually captured by the camera itself, although in 25  
 26 the case of recovering HDR information from a single negative the same technique 26  
 27 may be applied during the film-scanning phase. 27  
 28

29 By taking multiple exposures, each image in the sequence will have different 29  
 30 pixels properly exposed, and other pixels under- or overexposed. However, each 30  
 31 pixel will be properly exposed in one or more images in the sequence. It is therefore 31  
 32 possible and desirable to ignore very dark and very bright pixels in the subsequent 32  
 33 computations. 33

34 Under the assumption that the capturing device is perfectly linear, each exposure 34  
 35 may be brought into the same domain by dividing each pixel by the image’s expo- 35

1 sure time. From the recorded radiance values  $L_e$ , this effectively recovers irradiance  
 2 values  $E_e$  by factoring out the exposure duration.<sup>1</sup>

3 Once each image is in the same unit of measurement, corresponding pixels may  
 4 be averaged across exposures — excluding, of course, under- and overexposed pix-  
 5 els. The result is an HDR image.

6 In practice, cameras are not perfectly linear light measurement devices, objects  
 7 frequently do not remain still between individual exposures, and the camera is  
 8 rarely kept still. Thus, in practice this procedure needs to be refined to include cam-  
 9 era response curves, image alignment techniques, and ghost and lens flare removal.

10 Extracting a medium-dynamic-range radiance map from a single negative is rela-  
 11 tively straightforward because it does not require alignment of multiple frames,  
 12 and does not suffer from object displacement that may occur during the capture of  
 13 several exposures. It therefore serves as the basis for the techniques presented later  
 14 in this chapter.

### 16 4.3 FILM SCANNING

18 In the ideal case for creating an HDR image from multiple LDR exposures, the scene  
 19 or image should be completely static (e.g., an exposed and developed negative). We  
 20 assume that the response curve of the film is known. In addition, the LDR capture  
 21 device (such as an 8-bit/primary film scanner with known response curves) should  
 22 provide some means of exactly controlling the exposure during multiple captures.

23 Creating an HDR image under these conditions starts by taking scans with mul-  
 24 tiple exposures. In addition, the system response is inverted to get back to a linear  
 25 relation between scene radiances and pixel values. Each scanned image is multi-  
 26 plied by a calibration factor related to its exposure, and combined into an HDR  
 27 result. The only question is what weighting function to use in averaging together  
 28 the linear exposures. Of course, the lightest and darkest pixels at the limits of each  
 29 exposure should be excluded from consideration because these pixels are under- or  
 30 overexposed. But how should the pixels between be weighted?

31 .....  
 32  
 33 <sup>1</sup> The quantity captured by the camera is spectrally weighted radiance. As such, calling this quantity “radiance” is inappro-  
 34 priate. However, the spectral response curve is typically not the same as the CIE  $V(\lambda)$  curve, and therefore this quantity  
 35 also cannot be called “luminance” [18]. When the term *radiance* or *irradiance* is used, it should be understood that this  
 refers to spectrally weighted radiance and irradiance.

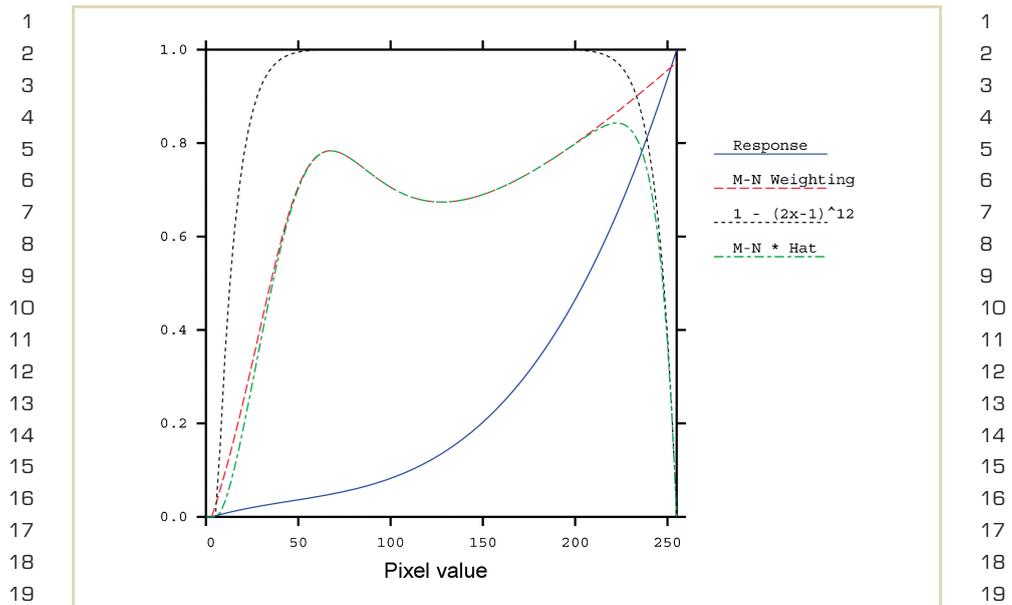


FIGURE 4.2 The inverted system response function (solid line) and recommended Mitsunaga–Nayar weighting function, multiplied by an additional hat function  $1 - (2x - 1)^{12}$  to devalue the extrema, which are often suspect.

Mann and Picard proposed a certainty/weighting function equal to the derivative of the system response curve for each color channel, using the argument that greater response sensitivity corresponds to greater certainty [5]. Debevec and Malik used a simple hat function based on the assumption that mid-range pixels are more reliable [18]. Mitsunaga and Nayar used signal theory to argue for multiplying Mann and Picard’s weight by the response output, in that larger values are less influenced by a constant noise floor [82]. Any of these methods will yield a satisfactory result, although the latter weighting function is better supported by signal theory. The Mitsunaga–Nayar weighting seems to work best when multiplied by a broad hat

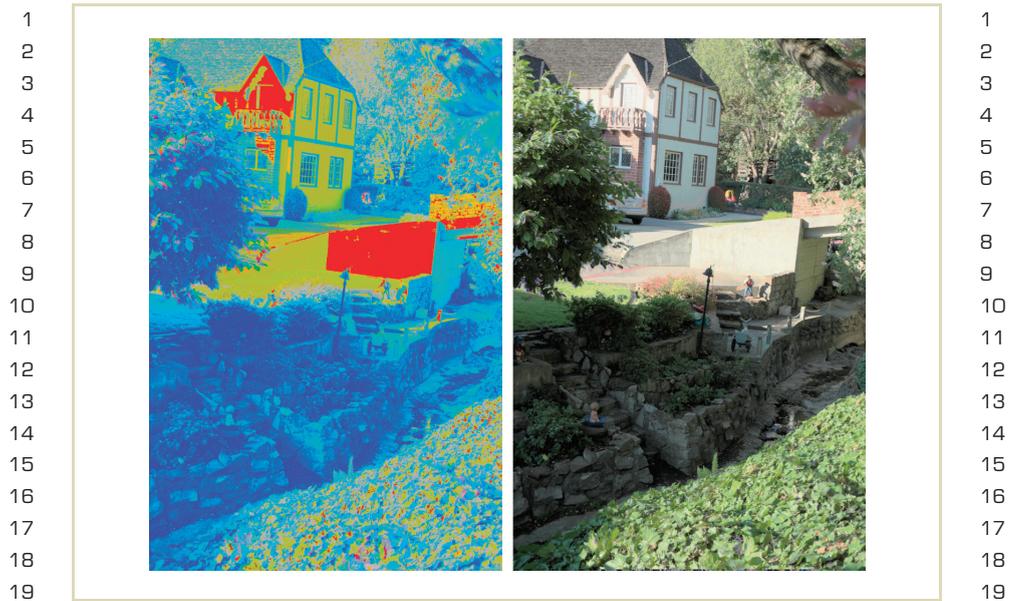


**FIGURE 4.3** Our example exposure sequence. Each image is separated by two f-stops (equal to a factor of 4, or  $0.6 \log_{10}$  units).

function, as shown in Figure 4.2. This avoids dubious pixels near the extremes, where gamut limitations and clamping may affect the output values unpredictably.

Figure 4.3 shows a sequence of perfectly aligned exposures. Figure 4.4 (left) shows the weighting used for each of the three contributing exposures, where blue is used for the longest exposure, green for the middle exposure, and red for the shortest exposure. As this figure shows, most pixels are a mixture of multiple exposures, with some pixels relying solely on the extremes of the exposure range. Figure 4.4 (right) shows the combined result, tone mapped using a histogram adjustment operator [142].

If the multiple exposures come not from multiple scans of a single negative but from multiple negatives or digital images, combining images may become problematic. First, the camera may shift slightly between exposures, which will result in some subtle (and possibly not-so-subtle) misalignments that will blur the re-



20  
21  
22  
23  
24  
25  
26

**FIGURE 4.4** The combined HDR result, tone mapped using the histogram adjustment operator (described in Section 7.2.8) in the right-hand image. The left-hand image shows contributing input image weights, where blue shows where the longer exposure dominates, green the middle, and red the shorter exposure. Most output pixels are an average of two or more input values, which reduces noise in the final result.

27  
28  
29  
30  
31  
32  
33  
34  
35

sults. Second, if the actual system response function is unknown the images must be aligned before this function can be estimated from the given exposures. Third, objects in the scene may shift slightly between frames or even make large movements, such as people walking in the scene as the photos are taken. Finally, flare in the camera lens may fog areas surrounding particularly bright image regions, which may not be noticeable in a standard LDR image. We will address each of these problems in turn and present some workarounds in the following sections.

29  
30  
31  
32  
33  
34  
35

## 4.4 IMAGE REGISTRATION AND ALIGNMENT

Although several techniques have been developed or suggested for image alignment and registration, most originating from the computer vision community, only two techniques to our knowledge address the specific problem of aligning differently exposed frames for the purpose of HDR image creation. The first technique, from Kang et al. [63], handles both camera movement and object movement in a scene, and is based on a variant of the Lucas and Kanade motion estimation technique [77]. In an off-line postprocessing step, for each pixel a motion vector is computed between successive frames. This motion vector is then refined with additional techniques, such as hierarchical homography (introduced by Kang et al.), to handle degenerate cases.

Once the motion of each pixel is determined, neighboring frames are warped and thus registered with one another. Then, the images are ready to be combined into an HDR radiance map. The advantage of this technique is that it compensates for fairly significant motion, and is suitable (for instance) for capturing HDR video by exposing successive frames by different amounts of time.

Although this method is suitable for significant motion, it relies on knowing the camera response function in advance. This presents a catch-22: alignment is needed to register samples to derive the camera response function, but the camera response function is needed to determine alignment. If the camera response is known or can be computed once and stored based on a set of perfectly aligned images, the catch-22 is solved.

A second alignment technique (described in following material) employs a *mean threshold bitmap* (MTB), which does not depend on the camera response function for proper alignment [141]. This technique is also about 10 times faster than the Kang et al. method, in that it performs its alignment operations on bitmaps rather than 8-bit grayscale images, and does not perform image warping or re-sampling. However, the MTB alignment algorithm does not address moving objects in the scene, and is not appropriate for arbitrary camera movements such as zooming and tilting. The method of Kang et al. may therefore be preferred in cases where arbitrary camera movement is expected. In the case of object motion, we recommend a simpler and more robust postprocessing technique in Section 4.7.

**4.5 THE MEAN THRESHOLD BITMAP ALIGNMENT TECHNIQUE**

In this section, we describe a method for the automatic alignment of HDR exposures [141].<sup>2</sup> Input to this exposure algorithm is a series of N 8-bit grayscale images, which may be approximated using only the green channel, or derived as follows from 24-bit sRGB with integer arithmetic.<sup>3</sup>

$$Y = (54 R + 183 G + 19 B)/256$$

One of the N images is arbitrarily selected as the reference image, and the output of the algorithm is a series of N-1 (x, y) integer offsets for each of the remaining images relative to this reference. These exposures may then be recombined efficiently into an HDR image using the camera response function, as described in Section 4.6.

The computation focuses on integer pixel offsets, because they can be used to quickly recombine the exposures without resampling. Empirical evidence suggests that handheld sequences do not require rotational alignment in about 90% of cases. Even in sequences in which there is some discernible rotation, the effect of a good translational alignment is to push the blurred pixels out to the edges, where they are less distracting to the viewer.

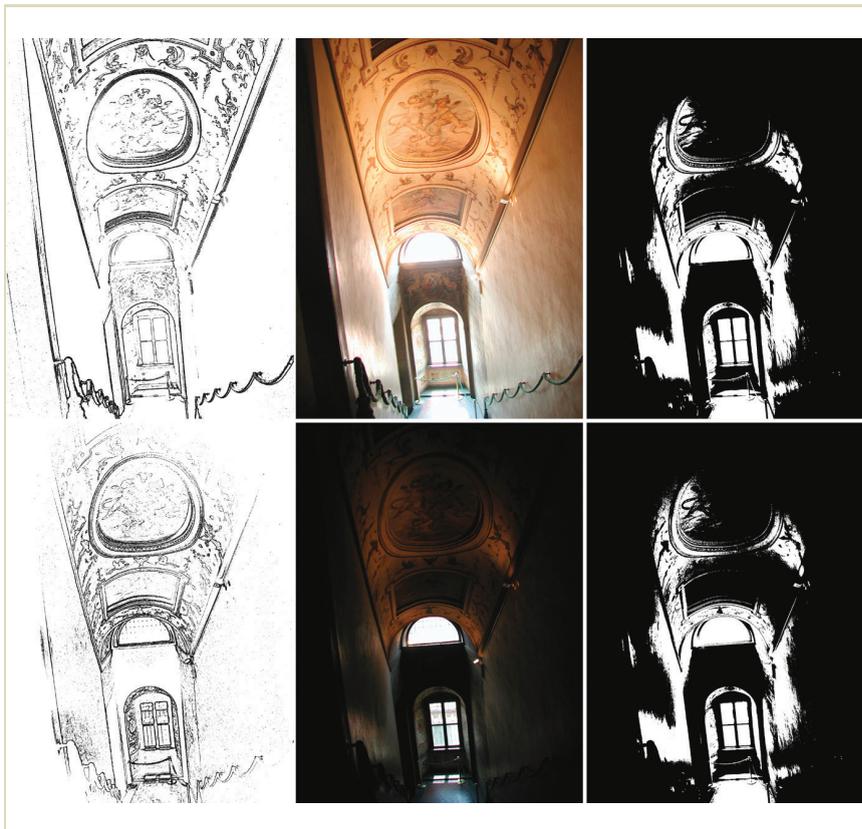
Conventional approaches to image alignment often fail when applied to images with large exposure variations. In particular, edge-detection filters are dependent on image exposure (as shown in the left side of Figure 4.5, where edges appear and disappear at different exposure levels). Edge-matching algorithms are therefore ill suited to the exposure alignment problem when the camera response is unknown. The MTB approach incorporates the following features.

- Alignment is done on bilevel images using fast bit-manipulation routines.
- The technique is insensitive to image exposure.
- For robustness, it includes noise filtering.

<sup>2</sup> Reprinted by permission of A.K. Peters, Ltd., from Greg Ward, "Fast, Robust Image Registration for Compositing High Dynamic Range Photographs from Hand-Held Exposures," *Journal of Graphics Tools*, 8(2):17-30, 2003.

<sup>3</sup> This is a close approximation of the computation of luminance as specified by ITU-R BT.709.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

**FIGURE 4.5** Two unaligned exposures (middle) and their corresponding edge bitmaps (left) and median threshold bitmaps (right). The edge bitmaps are not used precisely because of their tendency to shift dramatically from one exposure level to another. In contrast, the MTB is stable with respect to exposure.

1 The results of a typical alignment are discussed in Section 4.5.4. If we are to rely on 1  
 2 operations such as moving, multiplying, and subtracting pixels over an entire high- 2  
 3 resolution image, the algorithm is bound to be computationally expensive, unless 3  
 4 our operations are very fast. Bitmap images allow us to operate on 32 or 64 pixels 4  
 5 at a time using bitwise integer operations, which are very fast compared to byte- 5  
 6 wise arithmetic. We use a bitmap representation that facilitates image alignment 6  
 7 independent of exposure level, the a forementioned *median threshold bitmap*. The MTB 7  
 8 is defined as follows. 8  
 9

- 10 1 Determine the median 8-bit value from a low-resolution histogram over the 10
- 11 grayscale image pixels. 11
- 12 2 Create a bitmap image with 0s where the input pixels are less than or equal 12
- 13 to the median value, and 1s are where the pixels are greater. 13

14  
 15 Figure 4.5 shows two exposures of an Italian stairwell (middle), and their corre- 15  
 16 sponding edge maps (left) and MTBs (right). In contrast to the edge maps, the 16  
 17 MTBs are nearly identical for the two exposures. Taking the difference of these two 17  
 18 bitmaps with an exclusive-or (XOR) operator shows where the two images are mis- 18  
 19 aligned, and small adjustments in the  $x$  and  $y$  offsets yield predictable changes in 19  
 20 this difference due to object coherence. However, this is not the case for the edge 20  
 21 maps, which are noticeably different for the two exposures, even though we at- 21  
 22 tempted to compensate for the camera's nonlinearity with an approximate response 22  
 23 curve. Taking the difference of the two edge bitmaps would not give a good indi- 23  
 24 cation of where the edges are misaligned, and small changes in the  $x$  and  $y$  offsets 24  
 25 yield unpredictable results, making gradient search problematic. More sophisticated 25  
 26 methods of determining edge correspondence are necessary to use this information, 26  
 27 and we can avoid these and their associated computational costs with the MTB-based 27  
 28 technique. 28

29 The constancy of an MTB with respect to exposure is a very desirable property 29  
 30 for determining image alignment. For most HDR reconstruction algorithms, the 30  
 31 alignment step must be completed before the camera response can be determined, 31  
 32 in that the response function is derived from corresponding pixels in the differ- 32  
 33 ent exposures. An HDR alignment algorithm that depends on the camera response 33  
 34 function poses a catch-22 problem, as described earlier. By its nature, an MTB is 34  
 35 the same for any exposure within the usable range of the camera, regardless of the 35

1 response curve. As long as the camera’s response function is monotonic with respect to world radiance, the same scene will theoretically produce the same MTB at any exposure level. This is because the MTB partitions the pixels into two equal populations: one brighter and one darker than the scene’s median value. Because the median value does not change in a static scene, the derived bitmaps likewise do not change with exposure level.<sup>4</sup>

2  
3  
4  
5  
6

7 There may be certain exposure pairs that are either too light or too dark to use the median value as a threshold without suffering from noise, and for these we choose either the 17th or 83rd percentile as the threshold, respectively. Although the offset results are all relative to a designated reference exposure, we actually compute offsets between adjacent exposures and thus the same threshold may be applied to both images. Choosing percentiles other than the 50th (median) results in fewer pixels to compare, and this makes the solution less stable and thus we may choose to limit the maximum offset in certain cases. The behavior of percentile threshold bitmaps is otherwise the same as the MTB, including stability over different exposures. In the remainder of this section, when we refer to the properties and operations of MTBs, the same applies for other percentile threshold bitmaps.

8  
9  
10  
11  
12  
13  
14  
15  
16  
17

18 Once the threshold bitmaps corresponding to the two exposures have been computed, there are several ways to align them. One brute force approach is to test every offset within the allowed range, computing the XOR difference at each offset and taking the coordinate pair corresponding to the minimum difference. A more efficient approach might follow a gradient descent to a local minimum, computing only local bitmap differences between the starting offset (0,0) and the nearest minimum. We prefer a third method, based on an image pyramid that is as fast as gradient descent in most cases but more likely to find the global minimum within the allowed offset range.

19  
20  
21  
22  
23  
24  
25  
26

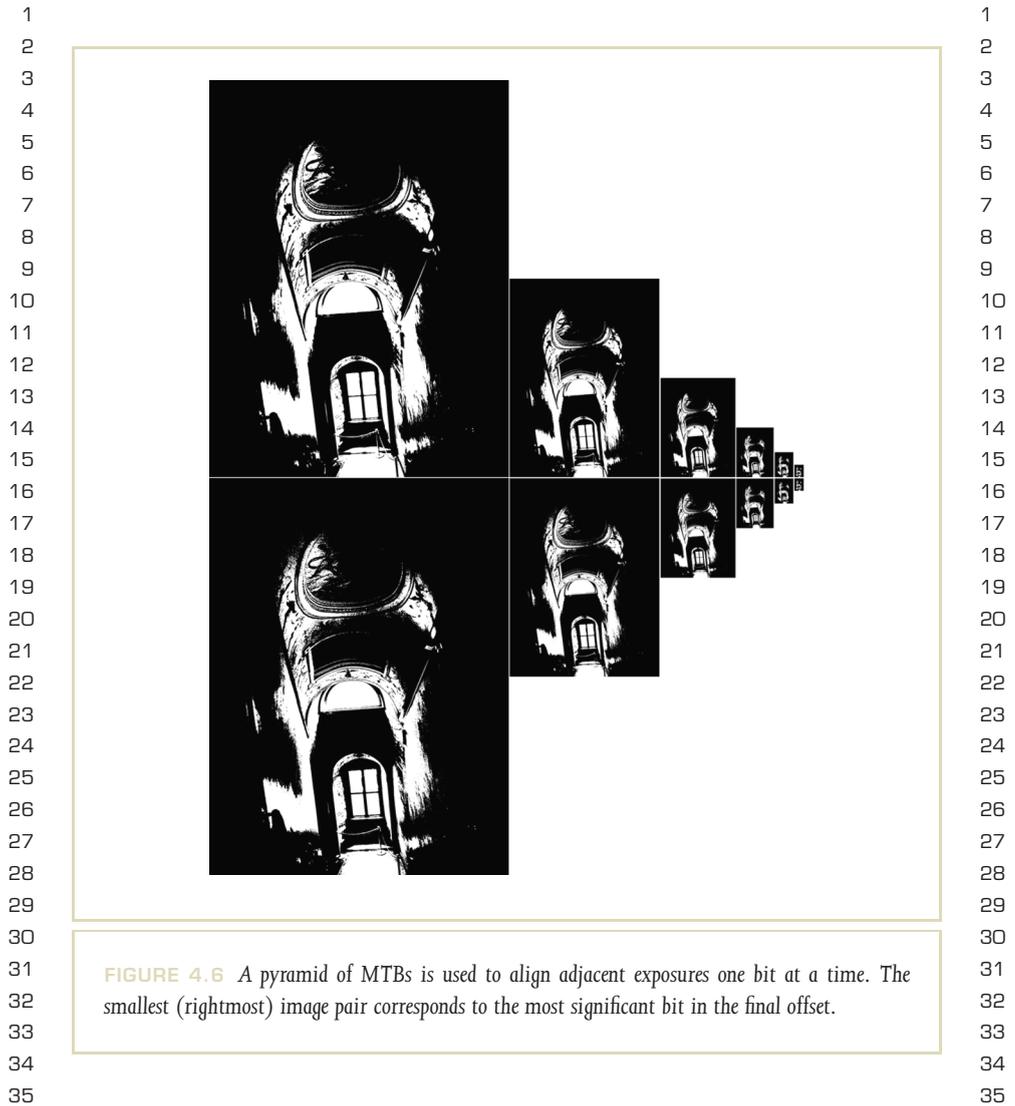
27 Multiscale techniques are well known in the computer vision and image-processing communities, and image pyramids are frequently used for registration and alignment. (See, for example, [127].) This technique starts by computing an image pyramid for each grayscale image exposure, with  $\log_2(\text{max\_offset})$  levels past the base resolution. The resulting MTBs are shown for two example exposures in Figure 4.6. For each smaller level in the pyramid, we take the previous grayscale

28  
29  
30  
31  
32  
33

34 <sup>4</sup> Technically, the median value could change with changing boundaries as the camera moves, but such small changes in the median are usually swamped by noise, which is removed by this algorithm.

35

4.5 THE MEAN THRESHOLD BITMAP ALIGNMENT TECHNIQUE



1 image and filter it down by a factor of two in each dimension, computing the MTB 1  
 2 from the grayscale result. The bitmaps themselves should *not* be subsampled, as the 2  
 3 result will be subtly different and could potentially cause the algorithm to fail. 3

4 To compute the overall offset for alignment, we start with the lowest-resolution 4  
 5 MTB pair and compute the minimum difference offset between them within a range 5  
 6 of  $\pm 1$  pixel in each dimension. At the next resolution level, we multiply this offset 6  
 7 by 2 (corresponding to the change in resolution) and compute the minimum dif- 7  
 8 ference offset within a  $\pm 1$  pixel range of this previous offset. This continues to the 8  
 9 highest-resolution (original) MTB, where we get the final offset result. Thus, each 9  
 10 level in the pyramid corresponds to a binary bit in the computed offset value. 10

11 At each level, we need to compare exactly nine candidate MTB offsets, and the 11  
 12 cost of this comparison is proportional to the size of the bitmaps. The total time 12  
 13 required for alignment is thus linear with respect to the original image resolution 13  
 14 and independent of the maximum offset, in that the registration step is linear in the 14  
 15 number of pixels, and the additional pixels in an image pyramid are determined by 15  
 16 the size of the source image and the (fixed) height of the pyramid. 16  
 17

#### 18 19 **4.5.1 THRESHOLD NOISE**

20 21 The algorithm just described works well in images that have a fairly bimodal bright- 21  
 22 ness distribution, but can run into trouble for exposures that have a large number of 22  
 23 pixels near the median value. In such cases, the noise in near-median pixels shows 23  
 24 up as noise in the MTB, which destabilizes the difference computations. 24  
 25

26 The inset in Figure 4.7 shows a close-up of the pixels in the dark stairwell ex- 26  
 27 posure MTB, which is representative of the type of noise seen in some images. 27  
 28 Computing the XOR difference between exposures with large areas such as these 28  
 29 yields noisy results that are unstable with respect to translation because the pix- 29  
 30 els themselves tend to move around in different exposures. Fortunately, there is a 30  
 31 straightforward solution to this problem. 31

32 Because this problem involves pixels whose values are close to the threshold, 32  
 33 these pixels can be excluded from our difference calculation with an *exclusion bitmap*. 33  
 34 The exclusion bitmap consists of 0s wherever the grayscale value is within some 34  
 35 specified distance of the threshold, and 1s elsewhere. The exclusion bitmap for the 35



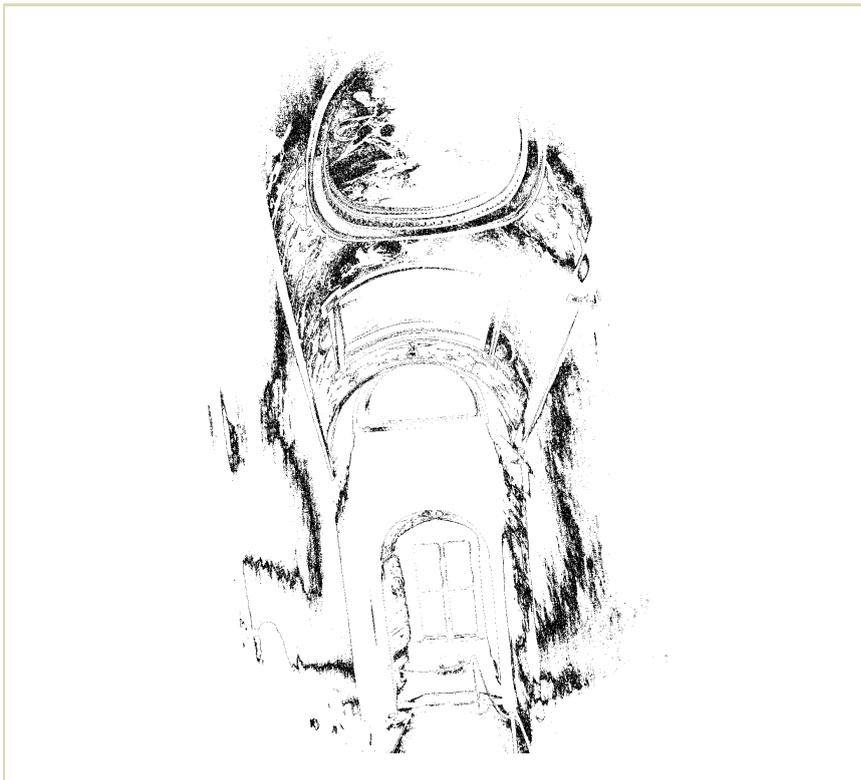
FIGURE 4.7 Close-up detail of noisy area of MTB in dark stairwell exposure (full resolution).

exposure in Figure 4.7 is shown in Figure 4.8, where all bits are zeroed for pixels within  $\pm 4$  of the median value.

We compute an exclusion bitmap for each exposure at each resolution level in the MTB pyramid, and then take the XOR difference result for each candidate offset, ANDing it with both offset exclusion bitmaps to compute the final difference.<sup>5</sup> The effect is to disregard differences that are less than the noise tolerance in our images.

<sup>5</sup> If we were to AND the exclusion bitmaps with the original MTBs before the XOR operation, we would inadvertently count disagreements about what was noise and what was not as actual pixel differences.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

**FIGURE 4.8** An exclusion bitmap, with zeroes (black) wherever pixels in our original image are within the noise tolerance of the median value.

This is illustrated in Figure 4.9, which shows the XOR difference of the unaligned exposures before and after applying the exclusion bitmaps. By removing those pixels that are close to the median, the least-reliable bit positions in the smooth gradients are cleared but the high-confidence pixels near strong boundaries (such as the edges



21  
22  
23  
24  
25

**FIGURE 4.9** The original XOR difference of the unaligned exposures (left), and with the two exclusion bitmaps ANDed into the result to reduce noise in the comparison (right).

26  
27  
28  
29  
30  
31

of the window and doorway) are preserved. Empirically, this optimization seems to be very effective in eliminating false minima in the offset search algorithm.

#### 4.5.2 OVERALL ALGORITHM

32 The full algorithm with the exclusion operator is given in the recursive C function  
33 (GetExpShift), shown in Figure 4.10. This function takes two exposure images,  
34 and determines how much to move the second exposure (*img2*) in *x* and  
35 *y* to align it with the first exposure (*img1*). The maximum number of bits in the

```

1  GetExpShift (const Image *img1, const Image *img2,
2              int shift_bits, int shift_ret[2])
3  {
4      int    min_err;
5      int    cur_shift[2];
6      Bitmap tb1, tb2;
7      Bitmap eb1, eb2;
8      int    i, j;
9      if (shift_bits > 0) {
10         Image sml_img1, sml_img2;
11         ImageShrink2(img1, &sml_img1);
12         ImageShrink2(img2, &sml_img2);
13         GetExpShift(&sml_img1, &sml_img2, shift_bits-1, cur_shift);
14         ImageFree(&sml_img1);
15         ImageFree(&sml_img2);
16         cur_shift[0] *= 2;
17         cur_shift[1] *= 2;
18     } else
19         cur_shift[0] = cur_shift[1] = 0;
20     ComputeBitmaps(img1, &tb1, &eb1);
21     ComputeBitmaps(img2, &tb2, &eb2);
22     min_err = img1->xres * img1->yres;
23     for (i = -1; i <= 1; i++)
24         for (j = -1; j <= 1; j++) {
25             int    xs = cur_shift[0] + i;
26             int    ys = cur_shift[1] + j;
27             Bitmap shifted_tb2;
28             Bitmap shifted_eb2;
29             Bitmap diff_b;
30             int    err;
31             BitmapNew(img1->xres, img1->yres, &shifted_tb2);
32             BitmapNew(img1->xres, img1->yres, &shifted_eb2);
33             BitmapNew(img1->xres, img1->yres, &diff_b);
34             BitmapShift(&tb2, xs, ys, &shifted_tb2);
35             BitmapShift(&eb2, xs, ys, &shifted_eb2);

```

**FIGURE 4.10** The GetExpShift algorithm.

```

1      BitmapXOR(&tb1, &shifted_tb2, &diff_b);
2      BitmapAND(&diff_b, &eb1, &diff_b);
3      BitmapAND(&diff_b, &shifted_eb2, &diff_b);
4      err = BitmapTotal(&diff_b);
5      if (err < min_err) {
6          shift_ret[0] = xs;
7          shift_ret[1] = ys;
8          min_err = err;
9      }
10     BitmapFree(&shifted_tb2);
11     BitmapFree(&shifted_eb2);
12     BitmapFree(&tb1); BitmapFree(&eb1);
13     BitmapFree(&tb2); BitmapFree(&eb2);
14 }

```

FIGURE 4.10 (Continued.)

final offsets is determined by the `shift_bits` parameter. The more important functions called by `GetExpShift` are as follows.

`ImageShrink2 (const Image *img, Image *img_ret)`: Sub-sample the image `img` by a factor of two in each dimension and put the result into a newly allocated image `img_ret`.

`ComputeBitmaps (const Image *img, Bitmap *tb, Bitmap *eb)`: Allocate and compute the threshold bitmap `tb` and the exclusion bitmap `eb` for the image `img`. (The threshold and tolerance to use are included in the `Image` struct.)

`BitmapShift (const Bitmap *bm, int xo, int yo, Bitmap *bm_ret)`: Shift a bitmap by  $(x_0, y_0)$  and put the result into the preallocated bitmap `bm_ret`, clearing exposed border areas to zero.

```

1   BitmapXOR (const Bitmap *bm1, const Bitmap *bm2, Bit- 1
2   map *bm_ret): Compute the XOR of bm1 and bm2 and put the result into 2
3   bm_ret. 3
4
5   BitmapTotal (const Bitmap *bm): Compute the sum of all 1 bits in 4
6   the bitmap. 5
7
8   Computing the alignment offset between two adjacent exposures is simply a 7
9   matter of calling the GetExpShift routine with the two image structs (img1 8
10  and img2), which contain their respective threshold and tolerance values. (The 9
11  threshold values must correspond to the same population percentiles in the two 10
12  exposures.) We also specify the maximum number of bits allowed in the returned 11
13  offset, shift_bits. The shift results computed and returned in shift_ret 12
14  will thus be restricted to a range of  $\pm 2^{\text{shift\_bits}}$ . 13
15
16  There is only one subtle point in this algorithm, which is what happens at the 14
17  image boundaries. Unless proper care is taken, non-zero bits may inadvertently 15
18  be shifted into the candidate image. These would then be counted as differences 16
19  in the two exposures, which would be a mistake. It is therefore crucial that the 17
20  BitmapShift function shifts 0s into the new image areas, so that applying the 18
21  shifted exclusion bitmap to the XOR difference will clear these exposed edge pixels 19
22  as well. This also explains why the maximum shift offset needs to be limited. In the 20
23  case of an unbounded maximum shift offset, the lowest-difference solution will also 21
24  have the least pixels in common between the two exposures (one exposure will end 22
25  up shifted completely off the other). In practice, we have found a shift_bits 23
26  limit of 6 ( $\pm 64$  pixels) to work fairly well most of the time. 24
27
28  4.5.3 EFFICIENCY CONSIDERATIONS 25
29
30  Clearly, the efficiency of the MTB alignment algorithm depends on the efficiency 26
31  of the bitmap operations, as nine shift tests with six whole-image bitmap oper- 27
32  ations apiece are performed. The BitmapXOR and BitmapAND operations are 28
33  easy enough to implement, as we simply apply bitwise operations on 32-bit or 64- 29
34  bit words, but the BitmapShift and BitmapTotal operators may not be as 30
35  obvious. 31
36
37  For the BitmapShift operator, any 2D shift in a bitmap image can be reduced 32
38  to a 1D shift in the underlying bits, accompanied by a clear operation on one or two 33
39  34
40  35

```

1 edges for the exposed borders. Implementing a 1D shift of a bit array requires at 1  
 2 most a left or right shift of  $B$  bits per word, with a reassignment of the underlying 2  
 3 word positions. Clearing the borders then requires clearing words where sequences 3  
 4 of 32 or 64 bits are contiguous, and partial clears of the remaining words. The 4  
 5 overall cost of this operator, although greater than the XOR or AND operators, is 5  
 6 still modest. This `BitmapShift` implementation includes an additional Boolean 6  
 7 parameter that turns off border clearing. This optimizes the shifting of the threshold 7  
 8 bitmaps, which have their borders cleared later by the exclusion bitmap, and thus 8  
 9 the `BitmapShift` operator does not need to clear them. 9

10 For the `BitmapTotal` operator, a table of 256 integers is computed corre- 10  
 11 sponding to the number of 1 bits in the binary values from 0 to 255 (i.e., 0, 1, 1, 11  
 12 2, 1, 2, 2, 3, 1, ..., 8). Each word of the bitmap can then be broken into chunks 12  
 13 (measured in bytes), and used to look up the corresponding bit counts from the 13  
 14 precomputed table. The bit counts are then summed to yield the correct total. This 14  
 15 results in a speedup of at least 8 times over counting individual bits, and may be 15  
 16 further accelerated by special-case checking for zero words, which occur frequently 16  
 17 in this application. 17

#### 18 4.5.4 RESULTS 18

19 Figure 4.11 shows the results of applying the MTB image alignment algorithm to 19  
 20 all five exposures of the Italian stairwell, with detailed close-ups showing before 20  
 21 and after alignment. The misalignment shown is typical of a handheld exposure 21  
 22 sequence, requiring translation of several pixels on average to bring the exposures 22  
 23 back atop each other. We have found that even tripod exposures sometimes need 23  
 24 minor adjustments of a few pixels for optimal results. 24  
 25 25

26 After applying this translational alignment algorithm to over 100 handheld ex- 26  
 27 posure sequences, a success rate of about 84% was found, with 10% giving un- 27  
 28 satisfactory results due to image rotation. About 3% failed due to excessive scene 28  
 29 motion — usually waves or ripples on water that happened to be near the threshold 29  
 30 value and moved between frames — and another 3% had too much high-frequency 30  
 31 content, which made the MTB correspondences unstable. Most of the rotation fail- 31  
 32 ures were mild, leaving at least a portion of the HDR image well aligned. Other 32  
 33 failures were more dramatic, throwing alignment off to the point where it was 33  
 34 better not to apply any translation at all. 34  
 35 35



**FIGURE 4.11** An HDR image composited from unaligned exposures (left) and detail (top center). Exposures aligned with the MTB algorithm yield a superior composite (right) with clear details (bottom center).

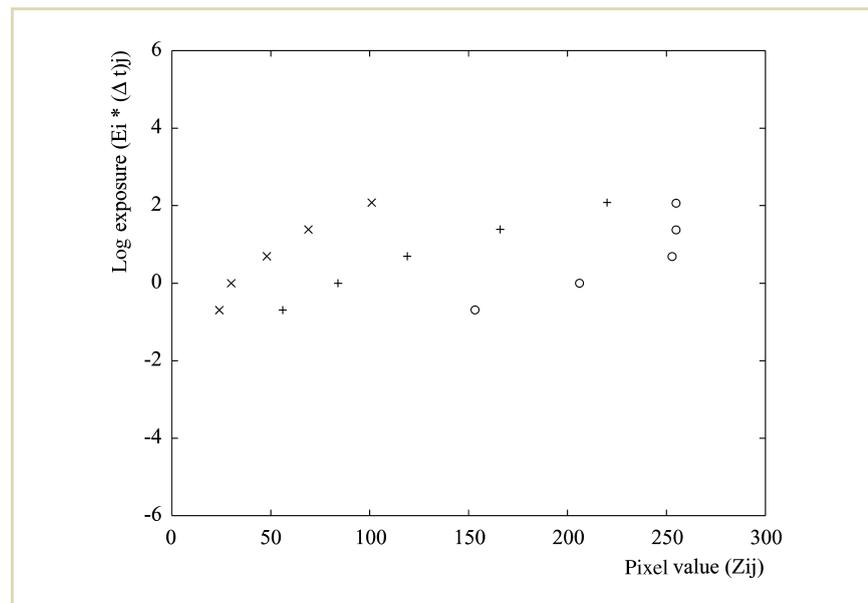
## 4.6 DERIVING THE CAMERA RESPONSE FUNCTION

Combining LDR exposures into an HDR image requires knowledge of the camera response function to linearize the data. In general, the response function is not provided by camera makers, who consider it part of their proprietary product differentiation. Assuming an sRGB response curve (as described in Chapter 2) is unwise, because most makers boost image contrast beyond the standard sRGB gamma to produce a livelier image. There is often some modification as well at the ends of the curves, to provide softer highlights and reduce noise visibility in shadows. However, as long as the response is not altered by the camera from one exposure to the next, it is possible to deduce this function given a proper image sequence.

1 **4.6.1 DEBEVEC AND MALIK TECHNIQUE**

2  
3 Debevec and Malik [18] demonstrated a simple and robust technique for deriving  
4 the camera response function from a series of aligned exposures, extending earlier  
5 work by Mann and Picard [5]. The essential idea is that by capturing different ex-  
6 posures of a static scene one is effectively sampling the camera response function at  
7 each pixel. This is best demonstrated graphically.

8  
9 Figure 4.12 shows three separate image positions sampled at five different expo-  
10 sures (Figure 4.13). The relative exposure ratios at each of the three positions are



31  
32  
33 **FIGURE 4.12** Plot of  $g(Z_{ij})$  from three pixels observed in five images, assuming unit radiance  
34 at each pixel. The images are shown in Figure 4.13.  
35



FIGURE 4.13 Three sample positions over five exposures shown in Figure 4.12.

given by the speed settings on the camera, and thus we know the shape of the response function at three different parts of the curve. However, we do not know how these three curve fragments fit together. Debevec and Malik resolved this problem using linear optimization to find a smooth curve that minimizes the mean-squared error over the derived response function. The objective function they use to derive the logarithmic response function  $g(Z_{ij})$  is as follows.

$$\mathcal{O} = \sum_{i=1}^N \sum_{j=1}^P \{w(Z_{ij}) [g(Z_{ij}) - \ln E_i - \ln \Delta t_j]\}^2 + \lambda \sum_{z=Z_{\min}+1}^{Z_{\max}-1} [w(z) g''(z)]^2$$

There,  $\Delta t_j$  is the exposure time for exposure  $j$ ,  $E_i$  is the film irradiance value at image position  $i$ , and  $Z_{ij}$  is the recorded pixel at position  $i$  and exposure  $j$ . The weighting function  $w(Z_{ij})$  is a simple hat function, as follows.

$$w(z) = \begin{cases} z - Z_{\min} & \text{for } z \leq \frac{1}{2}(Z_{\min} + Z_{\max}) \\ Z_{\max} - z & \text{for } z > \frac{1}{2}(Z_{\min} + Z_{\max}) \end{cases}$$

1 This equation is solved using singular-value decomposition to obtain an optimal  
 2 value for  $g(Z)$  for every possible value of  $Z$ , or 0 to 255 for an 8-bit image. Each  
 3 of the RGB channels is treated separately, yielding three independent response func-  
 4 tions. This assumes that interactions between the channels can be neglected. Al-  
 5 though this assumption is difficult to defend from what we know about camera  
 6 color transformations, it seems to work fairly well in practice.

7 Figure 4.14 shows the result of aligning the three curves from Figure 4.12, and  
 8 by applying the minimization technique to many image samples it is possible to  
 9 obtain a smooth response function for each channel.

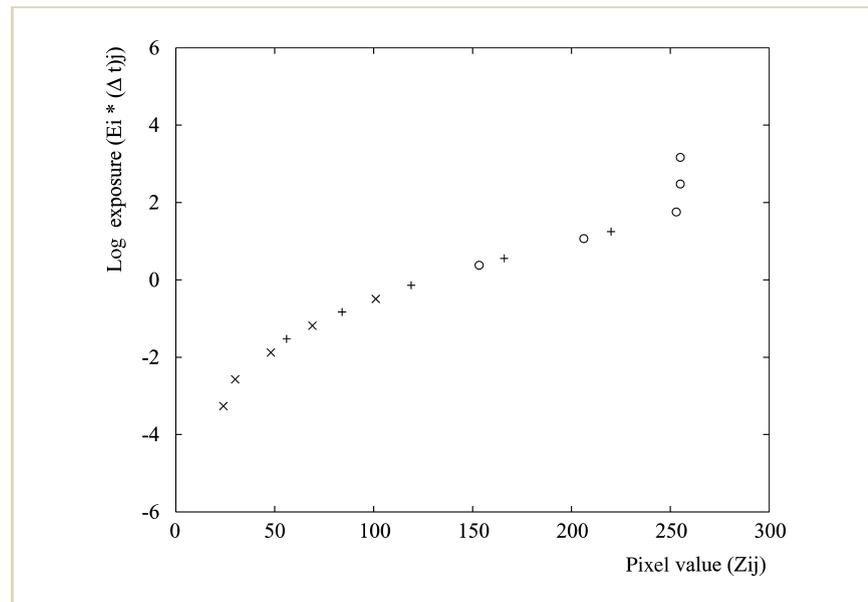


FIGURE 4.14 Normalized plot of  $g(Z_{ij})$  after determining pixel exposures.

1 **4.6.2 MITSUNAGA AND NAYAR TECHNIQUE** 1

2  
 3 Mitsunaga and Nayar presented a similar approach, in which they derive a poly- 3  
 4 nomial approximation to the response function [82] rather than the enumerated 4  
 5 table of Debevec and Malik. The chief advantage they cite in their technique is the 5  
 6 ability to resolve the exact exposure ratios in addition to the camera response func- 6  
 7 tion. This proves important for lower-cost consumer equipment whose aperture and 7  
 8 shutter speed may not be known exactly. Mitsunaga and Nayar define the following 8  
 9 N-dimensional polynomial for their response function: 9

$$10 \quad f(M) = \sum_{n=0}^N c_n M^n \quad 10$$

11  
 12 For consistency with Debevec and Malik, we suggest the following variable replace- 11  
 13 ments: 12  
 14

$$15 \quad N \rightarrow K \quad 15$$

$$16 \quad n \rightarrow k \quad 16$$

$$17 \quad M \rightarrow Z \quad 17$$

$$18 \quad Q \rightarrow P \quad 18$$

$$19 \quad q \rightarrow j \quad 19$$

20  
 21 The final response function is thus defined by the  $N + 1$  coefficients of this poly- 20  
 22 nomial,  $\{c_0, \dots, c_N\}$ . To determine these coefficients, they minimize the following 21  
 23 error function for a given candidate exposure ratio,  $R_{q,q+1}$  (the scale ratio between 22  
 24 exposure  $q$  and  $q + 1$ ): 23  
 25

$$26 \quad \varepsilon = \sum_{q=1}^{Q-1} \sum_{p=1}^P \left[ \sum_{n=0}^N c_n M_{p,q}^n - R_{q,q+1} \sum_{n=0}^N c_n M_{p,q+1}^n \right]^2 \quad 26$$

27  
 28 The minimum is found by determining where the partial derivatives with respect 27  
 29 to the polynomial coefficients are all zero (i.e., solving the following system of 28  
 30 31  
 32  
 33  
 34  
 35

1  $N + 1$  linear equations).

$$\frac{\partial \varepsilon}{\partial c_n} = 0$$

2  
3  
4  
5 As in previous methods, they only solve for the response up to some arbitrary scal-  
6 ing. By defining  $f(1) = 1$ , they reduce the dimensionality of their linear system by  
7 one coefficient, substituting

$$c_N = 1 - \sum_{n=0}^{N-1} c_n.$$

8  
9  
10 The final  $N \times N$  system can be written as follows.  
11  
12

$$\begin{bmatrix} \sum_{q=1}^{Q-1} \sum_{p=1}^P d_{p,q,0}(d_{p,q,0} - d_{p,q,N}) & \cdots & \sum_{q=1}^{Q-1} \sum_{p=1}^P d_{p,q,0}(d_{p,q,N-1} - d_{p,q,N}) \\ \cdots & \cdots & \cdots \\ \sum_{q=1}^{Q-1} \sum_{p=1}^P d_{p,q,N-1}(d_{p,q,0} - d_{p,q,N}) & \cdots & \sum_{q=1}^{Q-1} \sum_{p=1}^P d_{p,q,N-1}(d_{p,q,N-1} - d_{p,q,N}) \end{bmatrix} \times \begin{bmatrix} c_0 \\ \cdots \\ c_{N-1} \end{bmatrix} = \begin{bmatrix} - \sum_{q=1}^{Q-1} \sum_{p=1}^P d_{p,q,0} d_{p,q,N} \\ \cdots \\ - \sum_{q=1}^{Q-1} \sum_{p=1}^P d_{p,q,N-1} d_{p,q,N} \end{bmatrix}$$

13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25 Here,

$$d_{p,q,n} = M_{p,q}^n - R_{q,q+1} M_{p,q+1}^n.$$

26  
27  
28 The original Mitsunaga and Nayar formulation only considers adjacent exposures.  
29 In practice, the system is more stable if all exposure combinations are considered.  
30 The error function becomes a triple sum by including a sum over  $q' \neq q$  instead  
31 of just comparing  $q$  to  $q + 1$ . This then gets repeated in the sums of the combined  
32 system of equations, where  $d_{p,q,n}$  is replaced by  
33  
34

$$d_{p,q,q',n} = M_{p,q}^n - R_{q,q'} M_{p,q'}^n.$$

35

1 To compute the actual exposure ratios between images, Mitsunaga and Nayar apply  
 2 an interactive technique, where the previous system of equations is solved repeat-  
 3 edly, and between each solution the exposure ratios are updated using the following.

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

11 Iteration is complete when the polynomial is no longer changing significantly, as  
 12 follows.

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

15 This leaves just one final problem: What is the polynomial degree  $N$ ? The authors  
 16 recommend solving for every degree polynomial up to some maximum exponent  
 17 (e.g., 10), accepting the solution with the smallest error,  $\varepsilon$ . Fortunately, the solution  
 18 process proceeds quickly and this is not much of a burden. It is a good idea to  
 19 ensure that the same degree is selected for all color channels, and thus a combined  
 20  $\varepsilon$  function is preferable for this final test.

### 23 4.6.3 CHOOSING IMAGE SAMPLES FOR RESPONSE 24 RECOVERY

26 Each of the techniques described for camera response recovery requires a set of  
 27 intelligently selected samples from the exposure sequence. In principle, one could  
 28 use every pixel from every image, but this would only add to the computation time  
 29 while actually reducing stability in the solution due to misaligned and noisy data.  
 30 Once the exposures have been aligned, the following procedure for selecting sample  
 31 patches is recommended.

- 33 1 Sort the exposures from lightest to darkest.
- 34 2 Select an appropriate sample patch size and an optimal number of patches,  
 35 and initialize (clear) the patch list.

- 1     **3** Determine how many patches from the previous exposure are still valid for     1  
2     this one.     2
- 3     **4** Compute how many more patches are needed for this exposure. If none, go     3  
4     to the next exposure (Step 3).     4
- 5     **5** Search for valid patches using randomized rejection sampling. A valid patch     5  
6     is brighter than any of the previous exposure's patches, does not overlap any     6  
7     other patch, and possesses a low internal variance. It is also within the valid     7  
8     range for this exposure.     8
- 9     **6** Once we have found enough patches or given up due to an excess of rejected     9  
10    samples, we continue to the next exposure (Step 3).     10  
11

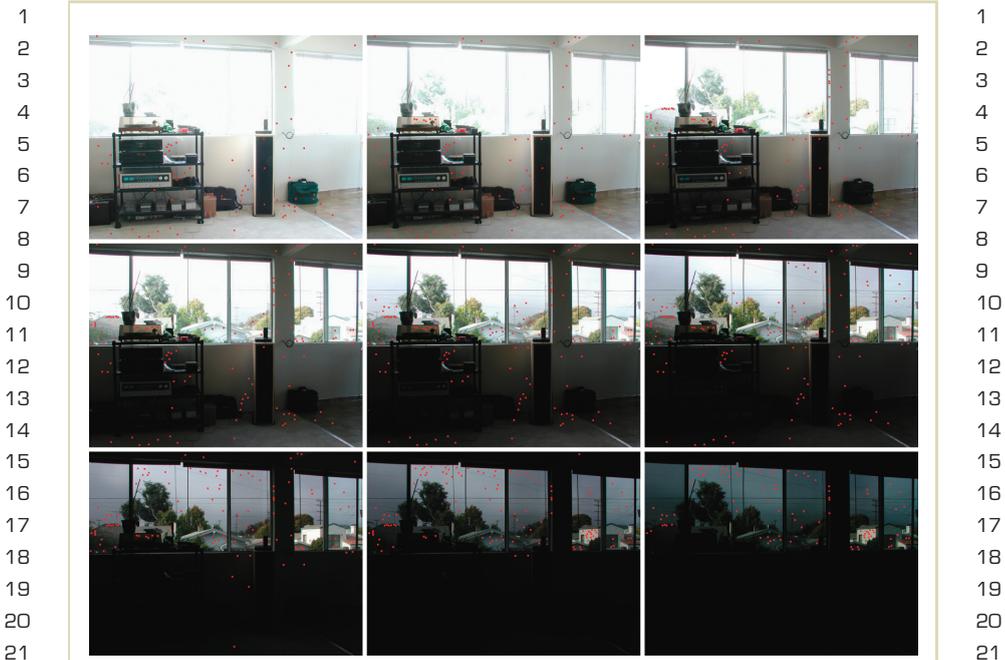
12 A target of 50 12-by-12 pixel patches per exposure seems to work well. In cases     12  
13 where the darker exposures do not use their full range, it becomes difficult to find     13  
14 new patches that are brighter than the previous exposure. In practice, this does not     14  
15 affect the result significantly, but it is important for this reason to place a limit on     15  
16 the rejection sampling process in Step 5, lest we go into an infinite loop.     16

17     Figure 4.15 shows an exposure sequence and the corresponding patch locations.     17  
18 Adjacent exposures have nearly the same patch samples, but no patch sample sur-     18  
19 vives in all exposures. This is due to the range restriction applied in Step 5 to avoid     19  
20 unreliable pixel values. Figure 4.16 shows a close-up of the middle exposure with     20  
21 the patches shown as boxes, demonstrating the low variance in the selected regions.     21  
22 By rejecting high-contrast areas, errors due to exposure misalignment and sensor     22  
23 noise are minimized.     23

24     Finally, Figure 4.17 shows the recovered response function for this sequence fit-     24  
25 ted with a third-order polynomial using Mitsunaga and Nayar's method, and com-     25  
26 pares it to the standard sRGB response function. The camera produces an artificially     26  
27 exaggerated contrast with deeper blacks on its LDR exposures. This type of response     27  
28 manipulation is fairly standard for consumer-grade cameras, and many professional     28  
29 SLRs as well.     29  
30

#### 31     **4.6.4 CAVEATS AND CALIBRATION**     31 32

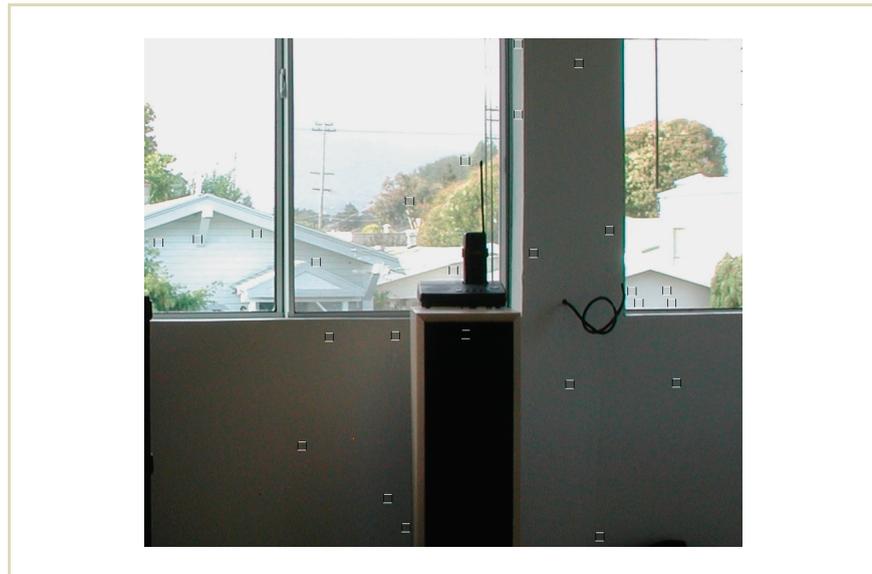
33  
34 To apply these techniques successfully, it helps to follow some additional guidelines,     34  
35 as follows.     35



**FIGURE 4.15** Red squares indicate size and location of patch samples in each exposure.

- Use aperture priority or manual exposure mode, so that only the exposure time is allowed to vary. This reduces problems associated with vignetting (light falloff toward the edge of the image).
- Fix the camera's white balance on a specific setting for the entire sequence, preferably daylight (i.e.,  $D_{65}$ ).
- If the camera offers an "optimized color and contrast" mode, switch it off. The more settings you can fix manually, the less likely the camera will be altering the response function between exposures. This applies particularly to automatic ISO/ASA and programmed exposure modes.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

**FIGURE 4.16** Close-up of central exposure showing selected patch regions.

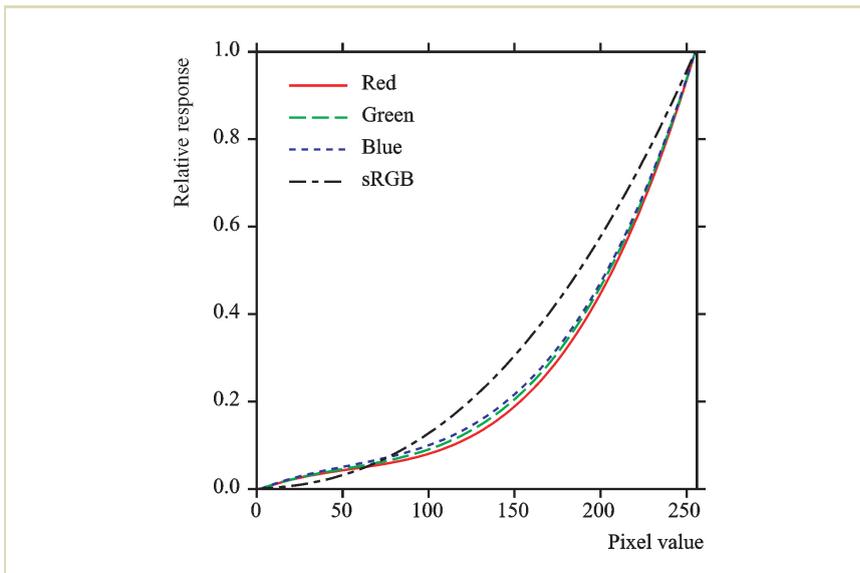
- Use a tripod if possible, and control your camera via a tether to a laptop computer if this option is available. The less touching of the camera during a sequence the fewer alignment problems you will experience.

In general, it works best to calibrate your camera's response one time, and then reuse this calibration for later exposure sequences. In this way, the scene and exposure sequence may be optimized for camera response recovery. For such a sequence, perform the following.

- Set the camera on a tripod and use a tether if available. Alignment may still be necessary between exposures if the camera is touched during the sequence, but the method described in the previous section will work far better with

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35



**FIGURE 4.17** Recovered red, green, and blue response functions for the image sequence shown in Figure 4.15.

a tripod than a handheld sequence in which the exposure is being changed manually.

- Choose a scene with large gray or white surfaces that provide continuous gradients for sampling. The closer your scene is to a neutral color the less likely color transforms in the camera will undermine the response recovery process.
- Choose a scene with very bright and very dark areas, and then take a long sequence of exposures separated by 1 EV (a factor of two in exposure time). The darkest exposure should have no RGB values greater than 200 or so, and the lightest exposure should have no RGB values less than 20 or so. Do not

1 include an excess of exposures beyond this range, as it will do nothing to 1  
 2 help with response recovery and may hurt. 2  
 3 • If you have access to a luminance meter, take a reading on a gray card or 3  
 4 uniform area in your scene to provide absolute response calibration. 4  
 5  
 6 Once a camera has been characterized in this way, it is possible to combine handheld 6  
 7 bracketed sequences that are too short to reliably recover the response function. 7  
 8  
 9 **4.7 GHOST REMOVAL** 9  
 10  
 11 Once the exposures are aligned with each other and the camera’s response curve is 11  
 12 determined, we may safely combine the images (as described in Section 4.3). How- 12  
 13 ever, if some person or object was moving during the image sequence acquisition 13  
 14 they may appear as “ghosts” in the combined result, due to their multiple locations. 14  
 15 The technique described by Kang et al. [63] attempts to address this problem dur- 15  
 16 ing alignment by warping pixels according to local content, but even if this can be 16  
 17 done correctly in the presence of people who change posture as well as position it 17  
 18 still leaves the problem of filling in holes that were obstructed in some views but 18  
 19 not in others. 19  
 20  
 21 A simpler approach is based on the observation that each exposure in the se- 21  
 22 quence is self-consistent, which means that we can simply choose one exposure or 22  
 23 another in specific regions to obtain a ghost-free result. The HDR capacity may be 23  
 24 lost within these selected regions, but as long as the ghosts are local and compact, 24  
 25 the overall image will still capture the full range of light. 25  
 26 Figure 4.18 shows an HDR image captured from a bracketed sequence of five 26  
 27 exposures, excerpted in the left-hand side of the figure. People walking in and out 27  
 28 of the temple result in a trail of ghosts appearing in the combined result. 28  
 29 Fortunately, it is relatively easy to detect motion of this type in exposure se- 29  
 30 quences. As the images are combined using the weighted average (described in 30  
 31 Section 4.3), the weighted variance can be computed simultaneously at each pixel, 31  
 32 shown in Figure 4.19. The weighted variance is defined as the weighted sum of 32  
 33 squares at each pixel over the square of the weighted average, the quantity minus 1. 33  
 34 (We compute these quantities separately for red, green, and blue channels and then 34  
 35 take the maximum at each pixel.) In addition to the moving people, some variance 35

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

**FIGURE 4.18** Five exposures combined into a single HDR image, where people moving through the scene have caused ghosting in the result.

is detected at high-contrast edges due to imperfections in the lens and sensor. These regions will usually be rejected by a minimum area constraint, but may cause false positives on small stationary objects.

At this point, the variance image could simply be thresholded and a single exposure selected to substitute for all high-variance pixels. However, this would incur significant artifacts. First, parts of moving objects whose pixels happened to correspond to the background locally would break apart. Second, and more seriously, choosing a single exposure for all high-variance pixels would result in excessive information loss, as problem pixels may be found in very different brightness regions.

The algorithm could be modified to pick the best exposure for each problem pixel, but this would create an even more serious breakup problem because different parts of the same object will be better exposed in different frames, where

4.7 GHOST REMOVAL

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35



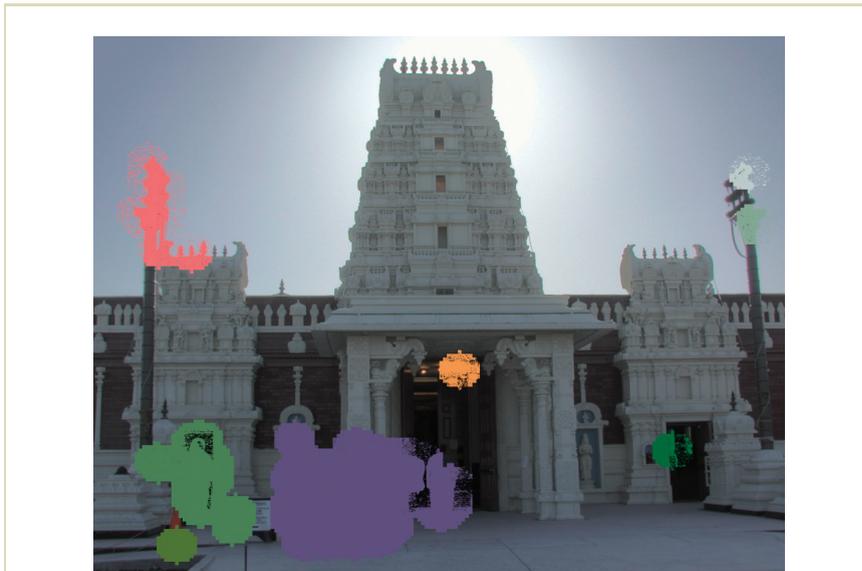
1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

**FIGURE 4.19** Variance computed at each pixel over our exposure sequence, showing where information is changing unexpectedly due to movement.

the object's position is also different. It is therefore important to isolate separable, high-variance regions and choose the best exposure for each. Such a segmentation is shown in Figure 4.20. This segmentation is computed as follows.

- 1 Reduce the variance image by a factor of 10 in each dimension to save computation.
- 2 Compute the threshold bitmap where local variance is greater than 0.18.
- 3 Smear the threshold bitmap around a radius of 3 pixels to cover edges and join adjacent ghost regions.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

**FIGURE 4.20** Segmented regions corresponding to isolated ghosts in our exposure sequence. Segment colors were randomly selected.

- 4 Compute a “background” bitmap segment from a union of contiguous (flood-filled) low-variance regions that cover at least 0.1% of the image.
- 5 Identify “ghost” bitmap segments as disjoint flood-filled regions within the background segment, with each ghost also covering at least 0.1% of the image.

In Figure 4.20, the tops of the two silhouetted poles were inadvertently picked, due to the high variance at their edges. In addition, a large region in which people were passing each other on the walk ended up as one segment (violet). Fortunately, this segment is contained in a similarly lighted region, and thus a single exposure is adequate to capture its dynamic range. The uneven edges of some regions indicate

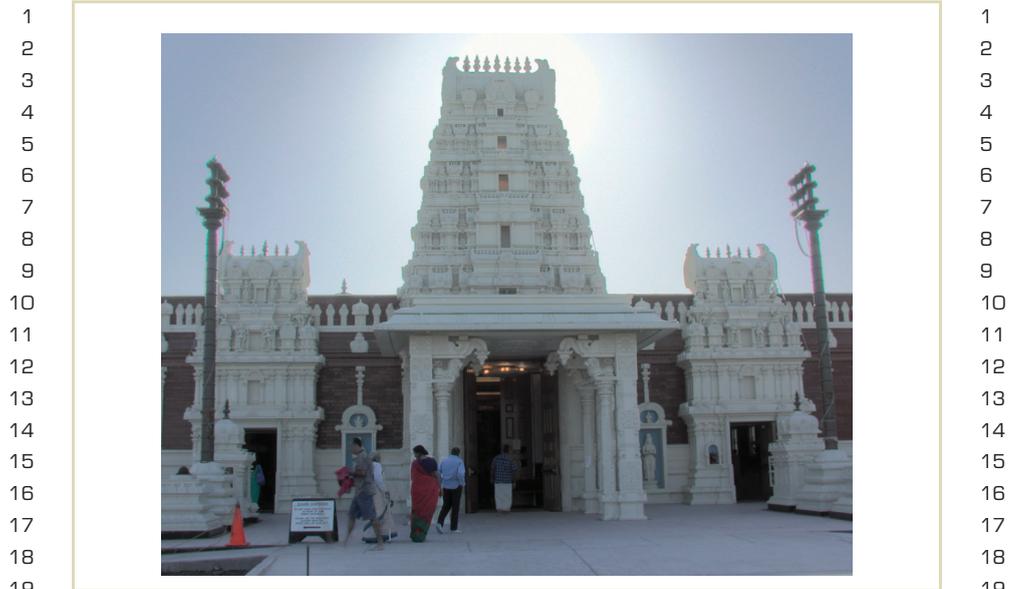


FIGURE 4.21 The combined HDR result with ghosts removed.

low-variance pixels, where we limit our ghost removal using linear interpolation as explained below.

To choose which exposure to use in which region, a histogram is generated from the floating-point values for each ghost segment. We then consider the largest value after ignoring the top 2% as outliers and choose the longest exposure that includes this 2% maximum within its valid range. We apply the corresponding exposure multiplier for each ghost segment then linearly interpolate between this exposure and the original HDR result using each pixel's variant as our mixing coefficient. This ensures that extremely low-variance pixels within an identified ghost segment are left unaltered. Figure 4.21 shows the combined result with ghosts removed. The final image is not perfect, and one man's bare foot has been summarily amputated,

1 but the overall result is an improvement, and this technique is quick and relatively  
2 straightforward.

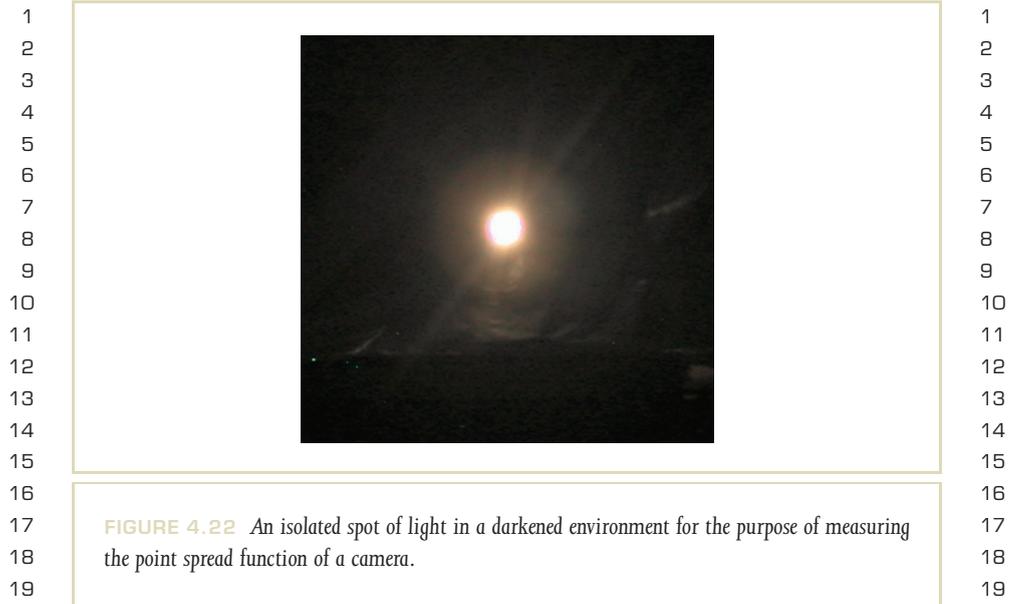
## 3 4 **4.8 LENS FLARE REMOVAL**

5  
6 After eliminating motion between exposures, there may still be artifacts present  
7 due to the camera's optics. Most digital cameras are equipped with optics that are  
8 consistent with the inherent limitations of 24-bit digital images. In other words,  
9 manufacturers generally do not expect more than two orders of magnitude to be  
10 captured in the final image, and thus certain parameters may be relaxed in the lens  
11 and sensor design relative to a 35-mm-film camera, for example. For an HDR cap-  
12 ture process, however, the limitations of the system's optics are more apparent, even  
13 in a well-made digital camera. Small issues such as the thickness and finish of the  
14 aperture vanes can make a big difference in the distribution of light on the sensor.  
15 The quality of coatings on the lenses and the darkness and geometry of the interior  
16 surrounding the sensor also come into play. Overall, there are many components  
17 that affect the scattering of light in an image, and it is difficult or impossible to  
18 arrive at a single set of measurements that characterize the system and its depen-  
19 dencies. Therefore, we prefer a dynamic solution to the lens flare problem, based  
20 only on the captured image.

21  
22 Because it is important to keep all of the optical properties of the camera consis-  
23 tent during HDR capture, normally only the shutter speed should be manipulated  
24 between exposures. Thus, the actual distribution of light on the sensor plane never  
25 varies, only the length of time the sensor is exposed to it. Therefore, any flare ef-  
26 fects present in one exposure are present to the same degree in all exposures and  
27 will sum consistently into our HDR result. For this reason, there is no need to work  
28 on individual exposures, as it would only serve to increase our computational bur-  
29 den. The camera's *point spread function* (PSF) is a physical measure of the system optics,  
30 and it may be characterized directly from the recorded radiances in an HDR image.

### 31 32 **4.8.1 THE POINT SPREAD FUNCTION**

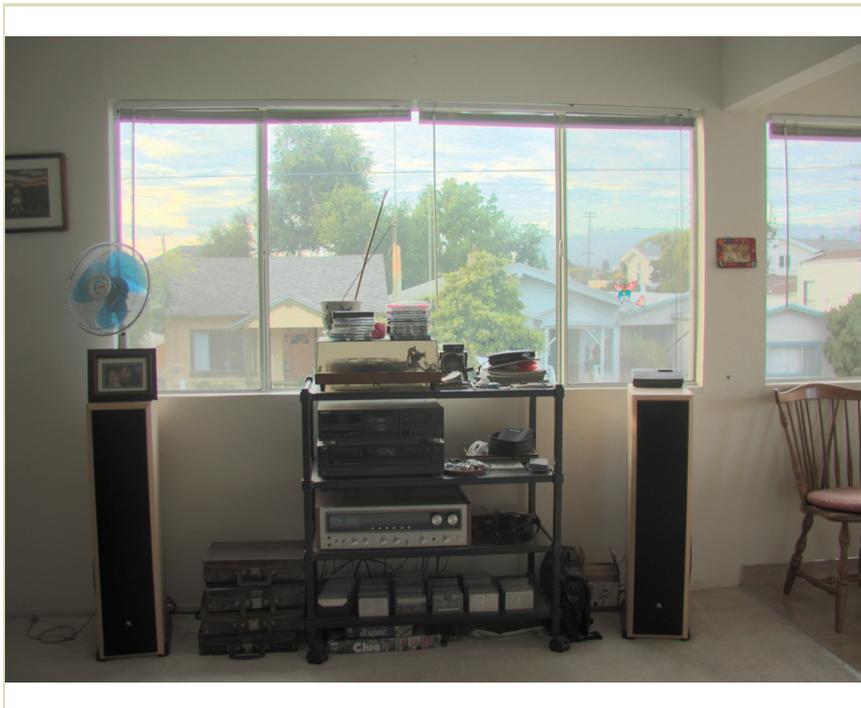
33  
34 The PSF as it is defined here is an idealized radially symmetric characterization of  
35 the light falloff surrounding a point of light in a perfectly dark surrounding. It



could be measured by making a pinhole in a piece of aluminum foil in front of a lightbulb in a box and photographing it in a completely dark environment, as shown in Figure 4.22. The edge of the hole ought to be perfectly sharp, but it generally is not. The spread of light around the hole corresponds to light scattered within the lens of the digital camera.

This photograph of the pinhole could then be used to correct the combined HDR result for other photographs made with precisely the same lens settings—zoom and aperture. However, this procedure is a lot to expect of even the most meticulous photographer, and because lens flare also depends strongly on dust and oils that come and go over time it is not practical to maintain a set of calibrated PSFs for any but the most critical applications.

However, there is a technique whereby the PSF may be approximated based on image content, as we will demonstrate with the HDR capture shown in Figure 4.23.



**FIGURE 4.23** Our input test image for estimating lens flare using the same aperture and zoom as in Figure 4.22 (tone mapped with the histogram equalization operator, described in Section 7.2.8).

Despite the fact that this image contains no localized bright spots, and hence no easily measured PSF, a reasonable estimate of lens flare may still be obtained.

It is assumed that in the image there exist some dark pixels near very bright (or “hot”) pixels. To the extent this is true, it will be possible to estimate the PSF

1 for the camera.<sup>6</sup> It is also assumed that the lens flare is radially symmetric. This is 1  
 2 admittedly a crude approximation, but required by the estimation procedure. Thus, 2  
 3 the goal is to find and remove the radially symmetric component of flare. Streaks 3  
 4 and other asymmetrical artifacts generated by the camera optics will remain. The 4  
 5 automatic flare removal consists of the following steps. 5  
 6

- 7 1 Compute two reduced-resolution HDR images: one in color and one in 7  
 8 grayscale. Call these  $I_{CR}$  and  $I_{GR}$ , respectively. 8
- 9 2 Identify “hot” pixels in  $I_{GR}$ , which are over some threshold. 9
- 10 3 Draw annuli around each hot pixel to compute a least squares approximation 10  
 11 to the PSF using the method described in the following section. 11
- 12 4 Apply the PSF to remove flare from the final HDR image. 12

13  
 14 Reducing the working resolution of our HDR image achieves a major speedup with- 14  
 15 out significantly impacting the quality of the results, in that flare tends to be a dis- 15  
 16 tributed phenomenon. A reduced image size of at most 128 pixels horizontally or 16  
 17 vertically is sufficient. The threshold setting for Step 2 is not particularly impor- 17  
 18 tant, but we have found a value of 1,000 times the minimum (reduced) pixel value 18  
 19 to work well for most images. Of course, a different threshold is advisable if the 19  
 20 minimum is zero. Steps 3 and 4 require some explanation, which we give in the 20  
 21 following subsections. 21  
 22

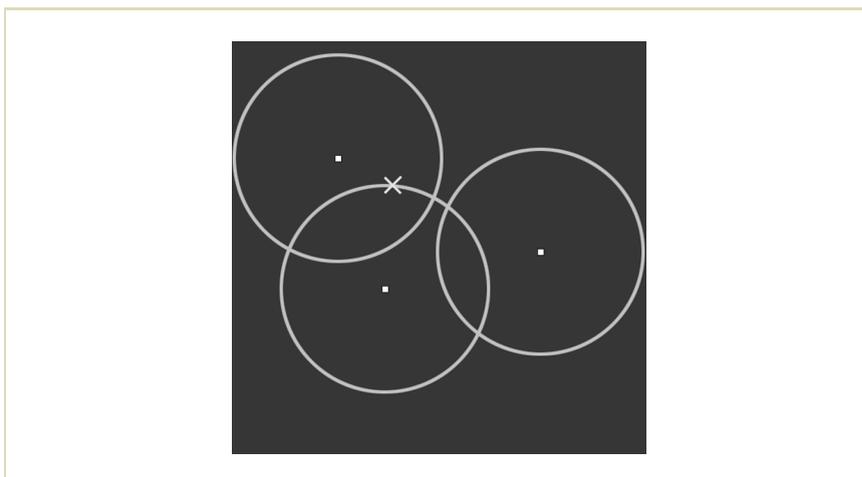
23 **4.8.2 ESTIMATING THE PSF** 23  
 24

25 The PSF defines how light falls off around bright points in the image.<sup>7</sup> To estimate 25  
 26 the PSF, the minimum pixel values around all “hot” pixels in the image are mea- 26  
 27 sured, thus arriving at a conservative estimate of the PSF. To do this, the potential 27  
 28 contributions of all hot pixels at a certain distance from the darker (non-hot) pixels 28  
 29 are summed to build up an estimate of the PSF from the corresponding minima, 29  
 30 radius by radius. For example, Figure 4.24 shows an image with exactly three of 30  
 31

32 ..... 32  
 33 <sup>6</sup> If this assumption is false, and there are no dark pixels near sources, lens flare will probably go unnoticed and there is 33  
 no need to remove it. 34

34 <sup>7</sup> In fact, it defines how light falls off around any point in the image, but only the bright points matter because the falloff is 34  
 35 so dramatic. 35

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22

**FIGURE 4.24** An example image with exactly three bright pixels, surrounded by circles showing where their PSF influences overlap. Each PSF radius will contain exactly one minimum, marked with an X in this example.

23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

these super-bright pixels. The same radius is drawn around all three pixels, creating three overlapping circles. If the PSF were known a priori, we could compute the contribution of these hot pixels at this distance by multiplying the PSF, which is a function of radius, by each hot pixel value. Conversely, dividing the darker pixels around each circle by the circle's center gives an upper bound of the PSF.

Furthermore, the PSF at this distance cannot be greater than the minimum of all darker-pixel/hot-pixel ratios. Where the circles overlap, the sum of hot pixel contributions should be considered. In fact, a convenient approach is to sum all three hot pixel values around each circle in another grayscale image. (This example shows three distinct circles, but in general there are many hot pixels adjacent to each other, which create a great deal of overlap in the contributing annuli.) The PSF upper bound at that radius will then equal the minimum ratio of the darker pixels in the circles to their hot pixel sums (the point marked with an X in the example). This

23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

- 1 technique extends directly to any number of hot pixels. The estimation procedure 1  
 2 is as follow. 2  
 3  
 4 **1** For each radius we wish to consider: 4  
 5 **a** Sum the hot pixel values into a radius range (annulus) in a separate 5  
 6 grayscale image. 6  
 7 **b** Find the minimum ratio of darker-pixel/hot-pixel sum for all annuli. 7  
 8 **2** If the minimum ratio is not less than the previous (smaller) radius, discard 8  
 9 it (because we assume the PSF is monotonically decreasing). 9  
 10 **3** For each minimum ratio pixel, identified for each sample radius, consider all 10  
 11 flare contributions to this pixel over the entire image as described below. 11

12  
 13 Once we have an estimate of the upper limit of the PSF at each radius, these mini- 13  
 14 mum pixels can be used to fit a third-degree polynomial,  $p(x)$ , using the reciprocal 14  
 15 of the input radius for  $x$ .<sup>8</sup> For each identified minimum pixel position with value 15  
 16  $P_i$ , we can write the following equation. 16

$$P_i = \sum_j P_j \left( C_0 + \frac{C_1}{r_{ij}} + \frac{C_2}{r_{ij}^2} + \frac{C_3}{r_{ij}^3} \right)$$

17  
 18 Here, the  $P_j$ s are the contributing pixel values over the rest of the image, and the 17  
 19  $r_{ij}$ s are the distances between the minimum pixel  $P_i$  and each contributing pixel 18  
 20 position. This equation can be rewritten as follows. 19  
 21

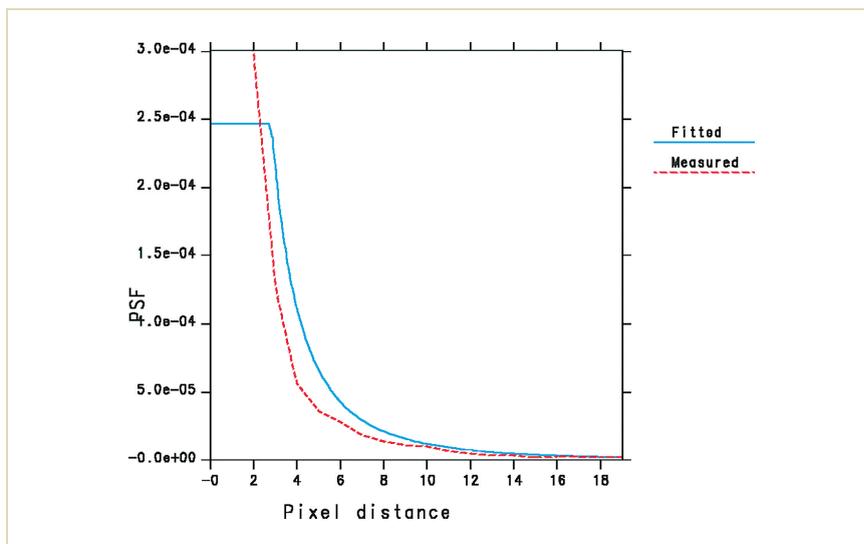
$$P_i = C_0 \sum_j P_j + C_1 \sum_j \frac{P_j}{r_{ij}} + C_2 \sum_j \frac{P_j}{r_{ij}^2} + C_3 \sum_j \frac{P_j}{r_{ij}^3}$$

22  
 23 The sums in this equation then become coefficients in a linear system in which the 23  
 24 four fitting parameters ( $C_0$  through  $C_3$ ) are the unknowns. As long as there are 24  
 25 more than four minimum pixel values,  $P_i$ , it should be possible to solve this as 25  
 26 an overdetermined system using standard least squares minimization. Heuristically, a 26  
 27 better solution may be obtained if we assign minimum and maximum permitted 27  
 28 values to the fitting parameters. 28  
 29  
 30  
 31  
 32  
 33

34 <sup>8</sup> The use of a third-degree polynomial and fitting to the reciprocal of distance are heuristic choices we have found to 34  
 35 produce good results at an economical cost. 35

1 values for the distance between pixels,  $r_{ij}$ . Anytime the actual distance is less than  
 2 the minimum radius (3 pixels in the reduced image), we use a distance of 3, instead. 2  
 3 Similarly, the distance is clamped to a maximum of half the image width. This 3  
 4 avoids stability problems and sensitivity to local features in our image. It also avoids 4  
 5 the possibly incorrect removal of flare too close to light sources. This is generally 5  
 6 impossible anyway, in that flare from the lens can be so great that the underlying 6  
 7 information is washed out. In such cases, no recovery can be made. This often 7  
 8 happens at bright source boundaries. 8

9 Figure 4.25 compares the point spread function measured in Figure 4.22 to our  
 10 estimate derived solely from the input image in Figure 4.23. Other than the artificial  
 11 plateau we imposed by constraining the minimum  $r_{ij}$  to 3 pixels, the two curves  
 12  
 13  
 14



15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

FIGURE 4.25 Comparison between directly measured PSF from Figure 4.22 and function fitted using image in Figure 4.23.

1 are a reasonable match. The fitted function shows a slightly greater flare than the 1  
 2 measured one, but this is explained by the fact that the measurement was based on 2  
 3 a spot near the center of the image. Optical flare becomes more pronounced as one 3  
 4 moves farther toward the edges of an image, especially in a wide-angle lens. Since 4  
 5 the fitting function was applied over the entire image, we would expect the globally 5  
 6 estimated PSF to be slightly greater than a PSF measured at the center. 6  
 7

### 8 4.8.3 REMOVING THE PSF 8

9  
 10 Given an estimate of the PSF, flare removal is straightforward. For each hot pixel in 10  
 11 the image, we subtract the PSF times this pixel value from its surroundings. Because 11  
 12 the neighborhood under consideration may extend all the way to the edge of the 12  
 13 image, this can be an expensive operation. Once again, working with a reduced 13  
 14 image lowers the computational cost to manageable levels. The steps for removing 14  
 15 the PSF are as follows. 15  
 16

- 17 1 Create a reduced-resolution flare image,  $\mathbf{F}_{CR}$ , and initialize it to black. 17
- 18 2 For each hot pixel in the reduced image  $\mathbf{I}_{CR}$ , multiply by the PSF and add the 18  
 19 product into  $\mathbf{F}_{CR}$ . 19
- 20 3 If the value of any pixel in  $\mathbf{F}_{CR}$  is larger than its corresponding pixel in  $\mathbf{I}_{CR}$ , 20  
 21 reduce the magnitude of  $\mathbf{F}_{CR}$  uniformly to compensate. 21
- 22 4 Upsample  $\mathbf{F}_{CR}$  using linear interpolation and subtract from the original HDR 22  
 23 image. 23

24  
 25 Step 3 ensures that no negative pixels are generated in the output and is necessary 25  
 26 because the fitting method does not guarantee the most conservative PSF. Dependent 26  
 27 on the interpolation and the local variance of the original pixels, we may still end 27  
 28 up with negative values during Step 4 and should truncate these where they occur. 28  
 29 An example result of automatic flare removal is shown in Figure 4.26, along with 29  
 30 the reduced resolution flare image generated during Step 2. 30  
 31

## 32 4.9 DIRECT CAPTURE OF HDR IMAGERY 32

33  
 34 With the possible exception of lens flare removal, the techniques explained in the 34  
 35 last section might be unnecessary if we had a digital sensor that could record the 35

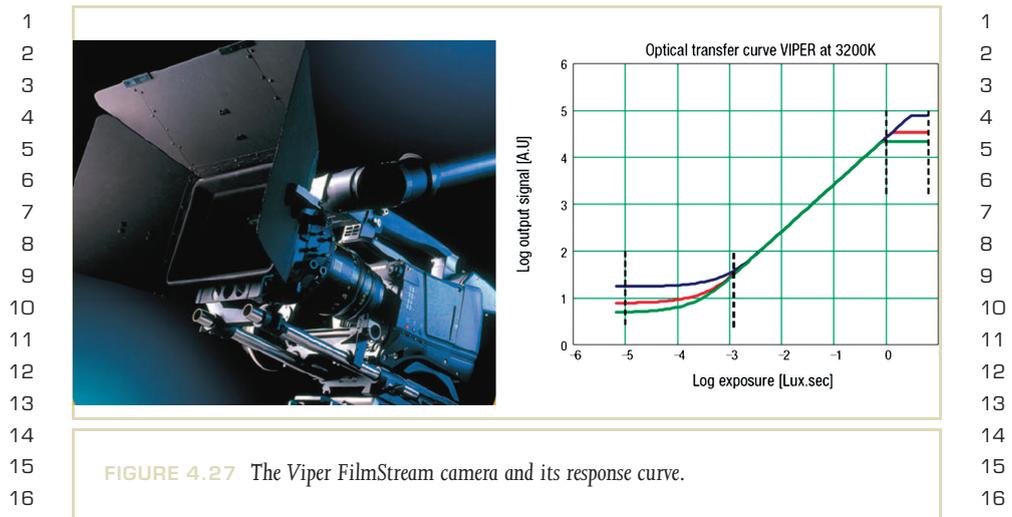


**FIGURE 4.26** An image of a rosette window, before flare removal (left) and after (center). The right-hand image shows the estimated PSF applied to hot pixels in the image.

full dynamic range of a scene in a single shot. In fact, such sensors are being actively developed, and some are even being marketed, but only a few integrated solutions are commercially available: the Autobrite cameras from SMaL Camera Technologies, the SpheroCam HDR panoramic camera from SpheronVR, and the Ladybug spherical camera from Point Grey Research. We will describe each of these systems briefly in this section.

#### 4.9.1 VIPER FILMSTREAM

Grass Valley, a division of Thomson, introduced the Viper FilmStream camera for digital cinematography in the Fall of 2002 ([www.thomsongrassvalley.com/products/cameras/viper/](http://www.thomsongrassvalley.com/products/cameras/viper/)). This is currently the top-end performer for digital capture, and it produces an enormous amount of data (up to 444 Mbytes/sec!). The camera contains three HDTV 1080i (1,920 × 1,080-resolution) CCD sensors (one each for red, green,



and blue) and records directly into a 10-bit/channel log format. The camera and its response functions are shown in Figure 4.27.

This chart shows that the Viper captures about three orders of magnitude, which is at least 10 times that of a standard digital video camera, and begins to rival film. The equipment is currently available for lease from Thomson.

#### 4.9.2 SMaL

SMaL Camera Technologies of Cambridge, Massachusetts ([www.smalcamera.com](http://www.smalcamera.com)), markets a low-cost VGA-resolution CMOS sensor (the IM-001 Series) which is capable of recording extended-range images at twice video rates (60 fps). Through its unique design, individual pixel sensitivities are adjusted so that the chip captures about twice the dynamic range (in log units) of a standard CCD or CMOS sensor, or about four orders of magnitude. They currently offer two products that incorporate their Autobrite (TM) technology, a credit-card-size still camera, and a video surveillance camera (a prototype is shown in Figure 1.9). They also market a “digital



**FIGURE 4.28** An HDR image of one of the authors, captured using a SMaL Ultra-Pocket digital camera directly to RGBE format. As we can see in false color, the dynamic range captured in this image is about three orders of magnitude.

imaging kit” to OEMs and system integrators who wish to incorporate the SMaL sensor in their products.

Figure 4.28 shows an image captured using the SMaL Ultra-Pocket camera. Due to its limited resolution ( $482 \times 642$ ), the SMaL sensor is not well suited to serious photographic applications, but this may change with the introduction of larger Autobrite arrays. Other aspects of chip performance, such as signal-to-noise ratio at each pixel and fill factor, may also affect the applicability of this technology.

### 4.9.3 PIXIM

Pixim of Mountain View, California ([www.pixim.com](http://www.pixim.com)), offers two  $720 \times 480$  CMOS image sensors that boast a 10-bit digital video output with a 95-dB signal-to-noise ratio, corresponding to roughly four orders of magnitude. These sensors grew out of the Programmable Digital Camera Project headed by Abbas El Gamal and Brian Wandell at Stanford University, and employ “multisampling” on picture elements to minimize noise and saturation. Pixels are grouped with independent analog-to-digital converters (ADCs), and sampled multiple times during each video frame.

1 Conversion at each sensor group stops when either the frame time is up or the 1  
 2 value nears saturation. In effect, each pixel group has its own electronic shutter 2  
 3 and dynamic exposure system. For additional processing, the sensor chip is paired 3  
 4 with a custom digital image processor, which handles video conversion and con- 4  
 5 trol. Pixim currently markets the sensors and development kits to OEMs, and it has 5  
 6 been picked up by a few security camera makers. Smartvue ([www.smartvue.com](http://www.smartvue.com)) has 6  
 7 based its S2 line of wireless surveillance cameras on the Pixim chip set, and Baxall 7  
 8 ([www.baxall.com](http://www.baxall.com)) recently introduced its Hyper-D camera. 8  
 9

#### 10 **4.9.4 SPHERONVR** 10

11 11  
 12 SpheronVR of Kaiserslautern, Germany ([www.spheron.com](http://www.spheron.com)), has what is undeniably the 12  
 13 highest-resolution and highest-performance HDR camera in existence; the Sphero- 13  
 14 Cam HDR. This device boasts the ability to capture full spherical panoramas at a 14  
 15 resolution of up to  $13,000 \times 5,300$  pixels, covering nearly eight orders of magni- 15  
 16 tude in dynamic range (a  $10^8:1$  contrast ratio). However, because they use a line- 16  
 17 scan CCD for their capture the process takes from 15 to 30 minutes to complete a 17  
 18 full 360-degree scan at this resolution. Lower resolutions and dynamic ranges will 18  
 19 scan faster, but one can never achieve a single-shot capture with a line-scan camera 19  
 20 because the device must mechanically pan over the scene for its exposure. Never- 20  
 21 theless, this is the system to beat for panoramic capture and critical image-based 21  
 22 lighting applications, and their deluxe package comes with an advanced software 22  
 23 suite as well. Figure 4.29 shows a SpheroCam HDR image captured in Napa Valley, 23  
 24 California, at a resolution of about  $3,000 \times 2,100$ . The dynamic range is 5.5 orders 24  
 25 of magnitude. 25  
 26

#### 27 **4.9.5 POINT GREY RESEARCH** 27

28 28  
 29 Point Grey Research, of Vancouver, Canada ([www.ptgrey.com](http://www.ptgrey.com)), recently came out with 29  
 30 an upgrade of the SDK for their LadyBug spherical video camera, which enables it 30  
 31 to capture six perspective HDR images in a single shot. Five 1/3-inch SVGA sensors 31  
 32 ( $1,024 \times 768$ ) look in a circle of horizontal directions to obtain a panorama, and 32  
 33 a sixth sensor looks straight up, yielding a final image that covers 75% of the full 33  
 34 sphere. Data are delivered in real time via a firewire cable to a tethered host com- 34  
 35 puter. Figure 4.30 shows an example HDR result with a resolution of  $3,600 \times 1,500$  35



**FIGURE 4.29** An HDR panorama captured by Spheron's SpheroCam HDR line-scan camera, tone mapped using a histogram adjustment operator.

and a dynamic range in excess of four orders of magnitude. Notably, there is no evidence of ghosting or smearing problems, which one would see if the image were multiply exposed.

#### 4.10 CONCLUSIONS

In the not-too-distant future digital still and motion picture photography may become exclusively HDR. After all, traditional film photography has provided medium-dynamic-range capture for nearly a century, and professionals expect and



**FIGURE 4.30** An HDR panorama captured in a single shot using the six CCD sensors of Point Grey Research's LadyBug camera system.

require this latitude during postproduction (i.e., printing). The current trend toward mixed reality in special effects is also driving the movie industry, which is increasingly digital, toward HDR. Advances in dynamic range will hit the professional markets first and slowly trickle into the semiprofessional price range over a period of years.

Unfortunately, consumers will continue to be limited to LDR digital cameras in the short term, as HDR equipment will be priced out of reach for some years to come. During this interim period, software algorithms such as those described in this chapter will be the most affordable way of obtaining and experimenting with HDR imagery, and applications will hopefully push the market forward.