

Multidimensional Image Enhancement from a Set of Unregistered and Differently Exposed Images

Ali Ajdari Rad, Laurence Meylan, Patrick Vandewalle and Sabine Süssstrunk
Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

ABSTRACT

If multiple images of a scene are available instead of a single image, we can use the additional information conveyed by the set of images to generate a higher quality image. This can be done along multiple dimensions. Super-resolution algorithms use a set of shifted and rotated low resolution images to create a high resolution image. High dynamic range imaging techniques combine images with different exposure times to generate an image with a higher dynamic range. In this paper, we present a novel method to combine both techniques and construct a high resolution, high dynamic range image from a set of shifted images with varying exposure times. We first estimate the camera response function, and convert each of the input images to an exposure invariant space. Next, we estimate the motion between the input images. Finally, we reconstruct a high resolution, high dynamic range image using an interpolation from the non-uniformly sampled pixels. Applications of such an approach can be found in various domains, such as surveillance cameras, consumer digital cameras, etc.

Keywords: Super-Resolution, High Dynamic Range, Camera Response Function

1. INTRODUCTION

Over the past years, an important trend in electronic imaging has been toward systems with higher resolution and higher dynamic range. A high resolution image is defined as an image with more resolving power, i.e., an image that contains more spatial information of the scene. According to this definition, an image that is obtained by upsampling and interpolating a low-resolution image does not have a higher resolution than its original. Although it has a greater number of pixels, it does not contain more details than the original image. One possibility to obtain a high resolution image is to use special sensors, which are often expensive. Another possibility is to use a software-based approach, called super-resolution. In such a setup, multiple low resolution images are combined to construct a higher resolution image.

Most super-resolution algorithms use low resolution images that differ only by small shifts, rotation or more complex motion, but not in exposure. They do not increase the dynamic range of the image. High dynamic range images can be constructed using a multiple exposure technique, which merges images captured with different exposure times into an image that contains the entire range of scene radiances. These two techniques are somehow similar but differ in the constraints that are imposed on the input images. For super-resolution, the input images must differ by small shifts or rotations but not in exposure time. For the multiple exposure technique, the input images must be registered (i.e., there should not be any motion between the images), but the exposure times must differ. In this article, we combine the two methods and construct a high quality image, having both a high resolution and a high dynamic range, from a set of shifted and rotated images taken with different exposure times. We then apply our method to movies and reconstruct a high resolution and high dynamic range movie from a set of frames taken with alternating exposure times.

This document is organized as follows: Section 2 presents the related work. Section 3 describes the proposed method. Experimental results are presented in Section 4, including data gathering methods, and development approaches. Section 5 concludes this article.

Send correspondence to Ajdari Rad A., E-mail: ali.ajdarirad@epfl.ch, Telephone: +41 21 69 32612.

This paper is reproducible. More results and the Matlab code used in this paper can be found at: http://lcavwww.epfl.ch/reproducible_research.

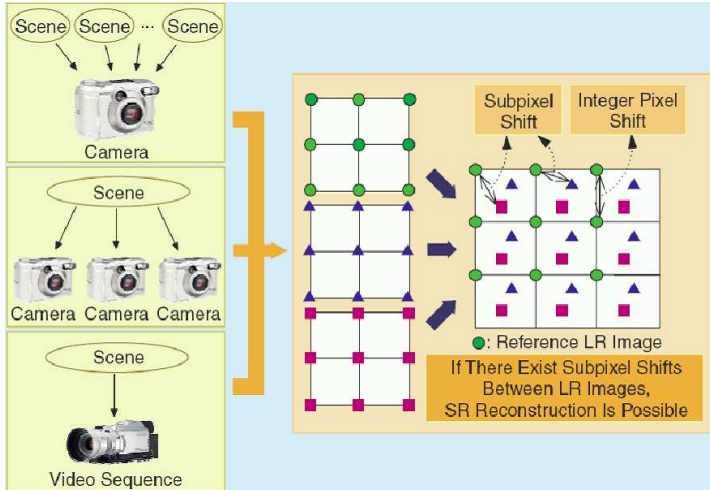


Figure 1. Principle of super-resolution algorithms.¹

2. BACKGROUND

In this section, we review the background work on super-resolution and construction of high dynamic range images. Then, we describe a few methods that combine both approaches.

2.1. Super-Resolution

An image with high spatial resolution can be obtained by cameras with high resolving power. The problems with these cameras are their high cost, their large size and the fact that they require high quality optics. Software-based approaches offer solutions to increase the spatial resolution that do not require the use of expensive hardware.

The common principle of super-resolution algorithms is to interpolate a high resolution image from multiple low resolution images that differ by small motions, typically sub-pixel shifts and rotations (see Figure 1).¹ The input to a super-resolution algorithm is a set of (shifted and rotated) low resolution images. This input set can be captured in several acquisitions using a single camera, multiple acquisitions with several cameras, or it can be extracted from a video as a sequence of frames. Note that the scene objects should remain fixed (or almost fixed) during the image acquisition. Also, the light conditions are often considered unchanged. Thus in practice, the only acceptable differences between the input images are small shifts and rotations.

A super-resolution algorithm computes these shifts and rotations (called registration or motion estimation) and constructs a higher resolution image from the registered low resolution input images. For a comprehensive review of super-resolution imaging, we refer to the papers by Farsiu et al.² and Park et al.¹ The existing image registration methods can be divided into frequency domain and spatial domain methods. In this paper, we use the frequency domain algorithm developed by Vandewalle et al.³ This algorithm is fast and accurate in comparison with other algorithms, such as the frequency domain algorithms by Marcel et al.,⁴ Lucchese and Cortelazzo,⁵ or the spatial domain algorithm by Keren et al.⁶ As opposed to most other methods, it can also deal with aliased input images. Another approach taking aliasing into account is the algorithm by Robinson et al.⁷

2.2. Construction of a High Dynamic Range Image

Standard capturing devices can only capture a limited dynamic range. In the case of a high dynamic range scene whose luminance ratio exceeds the dynamic range of the camera sensor, the captured image can only include tonal details corresponding to a limited range of scene radiances. Either details in the shadow are visible and the

highlights are clipped, or the details in the shadow are not visible due to noise (insufficient signal to noise ratio). Multiple exposure techniques have been developed to overcome the limitations of standard capturing devices and record the entire range of the scene radiances even if the scene’s dynamic range is larger than that of the device.⁸⁻¹¹ The principle is to capture the scene several times with different exposure settings and reconstruct a high dynamic range image that contains all the details even in bright and dark areas. The first step of these methods is to calculate the camera response function (CRF). Its inverse is then applied to the captured images in order to make them linear with respect to scene radiances. Next, a factor is applied to the pixel values to compensate for the different exposure times. Noisy and saturated pixels are discarded and the high dynamic range image is constructed by taking a weighted average of the valid pixel values.¹²⁻¹⁵

2.3. Combination of Super-Resolution and Multiple Exposure Technique

The combination of super resolution and multiple exposure techniques has not been addressed by many other authors. A solution using special hardware was proposed by Narasimhan and Nayar¹⁶ and Gunturk and Gevrekci,¹⁷ but to our knowledge software solutions do not exist. The related problem of registration of differently exposed images has been addressed by Ward,¹⁸ but it was not combined with super-resolution. Ward’s idea is to apply a thresholding on the pixel values to accelerate operations and to avoid problems with the varying exposure levels. We will take inspiration from this technique and adapt it to our purpose.

3. PROPOSED METHOD

An overview of our algorithm is given in Figure 2. The first step after capturing the images with different exposure times is to estimate the camera response function for each camera. We then apply the inverse camera response function on the input images and divide by the exposure time. At this stage, the images differ only in their extreme pixel values, i.e. the bright and dark pixels. This is due to the fact that in some images, some pixels are saturated or noisy and thus compensating for the exposure time does not recover the real data. Consequently, the second step consists in finding a common area that contains only pixels with meaningful data. This area is determined by a binary mask called the exposure invariant mask (see Figure 4), where black (0) represents an invalid pixel (noisy or saturated) and white (1) represents a valid pixel. The exposure invariant images are obtained by only considering image pixels where the mask is 1. Finally, we use the exposure invariant images to register the input images by applying the algorithm described by Vandewalle et al.³ to estimate the motion (shift and rotation).

Once the shift and rotation parameters are estimated, the last step is to perform an interpolation between the pixels of the input images, whose value has been changed to compensate for the different exposure times. To display the result, an appropriate tone mapping algorithm must be performed in order to take advantage of the high dynamic range of the computed image. As this is not the topic of this article, we will not address this issue here. All our figures were tone-mapped with a simple gamma correction.

3.1. Estimating the Camera Response Function

If the camera response function (CRF) of the capturing device is unknown, it has to be estimated first. Our estimation method is based on the algorithm presented by Debevec et al.⁸ Since this method assumes that the scene is unchanged and there is no shift or rotation between the input images (the images differ only in exposure time), the original algorithm could not be directly applied. Inspired by the algorithm proposed by Ward,¹⁸ we present a simple modification of Debevec et al.’s method that helps make the CRF estimation robust against shifts and rotations.

Suppose that we have three images of a scene with different exposure times and call them the normal, dark and bright image. Let us also assume that the imaging system is ideal and there is no noise in the captured images. The darkest area in the dark image corresponds to the darkest area in the normal and the bright image. This means that the histograms of these three images have the same shape and thus are invariant with respect to shift and rotation. Therefore, instead of applying Debevec et al.’s method on the “spatial” image, we apply

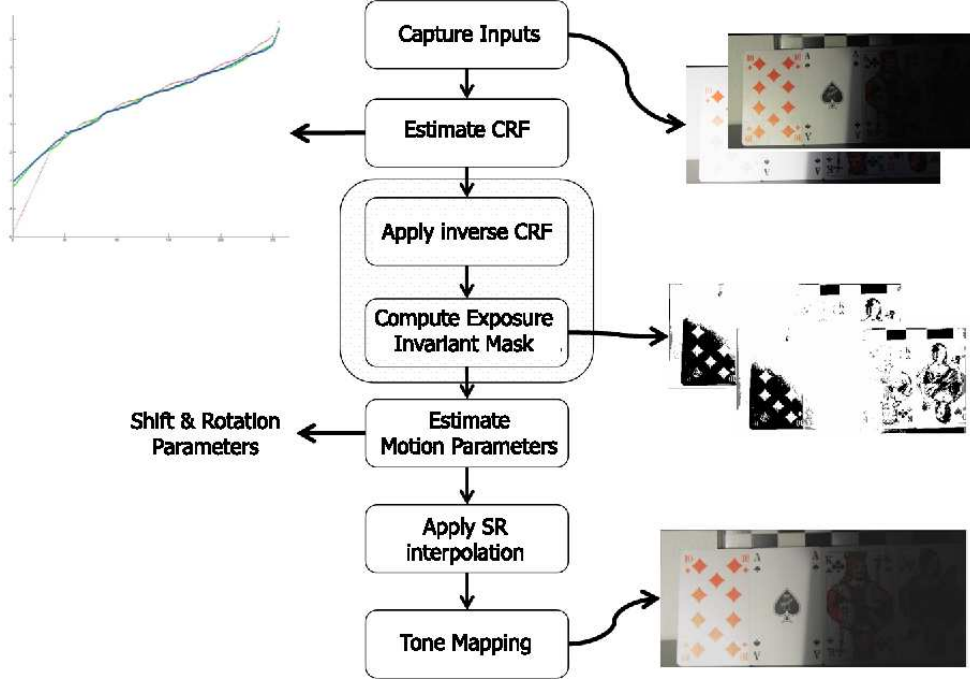


Figure 2. Block diagram of the proposed method.

it on the histogram of the image. This simple change gives significant improvements in the CRF estimation (see Section 4.2). The camera response function is estimated for each of the color channels, as well as for the grayscale version of the color image.

The input to our algorithm is a number of pictures $I_j(p)$ (usually four images) taken from the same scene with different known exposure times and unknown shifts and rotations. Image j is taken with a given exposure time t_j . The index p is used to indicate a specific pixel in the image I_j . We assume that the scene is static and that the acquisition is fast enough such that lighting changes between the images can be safely ignored. After sorting the pixels based on their values, we obtain the histogram of $I_j(p)$ that we write as $I'_j(p)$. Based on Debevec et al.'s algorithm, the relation between pixels and irradiances can be written as follows:

$$g(I'_j(p)) = \ln(f^{-1}(I'_j(p))) = \ln(E'_j(p) \cdot t_j) = \ln(E'_j(p)) + \ln(t_j). \quad (1)$$

In this set of equations, the pixel values $I'_j(p)$ and the exposure time t_j are known. The unknowns are the irradiances $E'_j(p)$, as well as the function g , which is the logarithm of the inverse camera response function. It is also assumed that g is smooth and monotonic. Debevec et al.'s method recovers simultaneously the camera response function g and the irradiances $E'_j(p)$ that best satisfy the set of equations arising from (1) in a least squares error sense. Let I_{min} and I_{max} be the lowest and highest pixel values (integers), P the number of pixels and N the number of images. The problem can then be formulated as a minimization of the following quadratic objective function:

$$\phi = \sum_{j=1}^N \sum_{p=1}^P w(I'_j(p))^2 \cdot (g(I'_j(p)) - \ln(E'_j(p)) - \ln(t_j))^2 + \lambda \sum_{x=I_{min}+1}^{I_{max}-1} (w(x)g''(x))^2, \quad (2)$$

where

$$w(x) = \begin{cases} x - I_{min} & \text{for } x \leq \frac{1}{2}(I_{min} + I_{max}) \\ I_{max} - x & \text{for } x \geq \frac{1}{2}(I_{min} + I_{max}). \end{cases} \quad (3)$$

The variable x denotes a pixel value in the range $[I_{min}, I_{max}]$, and g'' is the second derivative of the function g .

The first term ensures that the solution satisfies the set of equations arising from (1) in a least squares sense. The second term is a smoothness term on the sum of squared values of the second derivative of g . Because it is quadratic in $E'_j(p)$ and $g(x)$, minimizing ϕ is a straightforward linear least squares problem. The overdetermined system of linear equations is robustly solved using a singular value decomposition.⁸

3.2. Exposure Invariant Transform

Since our image registration method is sensitive to exposure variation between images, we need to extract an exposure invariant component from each image. This is done using the following equation:

$$\ln(E_j(p)) = g(I_j(p)) - \ln(t_j), \quad (4)$$

where $E_j(p)$ is the images, whose pixel values do not depend on the exposure time.

Applying the inverse of the CRF to the images and dividing them by the exposure times make them theoretically similar. However, they are still different in the saturated parts and in pixels with values near zero, which are not valid. These invalid pixels should not be used for the registration.

Our method transforms an input RGB image into grayscale values by extracting the luminance defined by a weighting average of the three color channels. To obtain the exposure invariant image, we apply a binary thresholding operation, where we discard the 5% darkest and brightest pixels. Note that the 5% darkest pixels in the brighter image and the 5% brightest pixels in the darker image are not generally saturated or noisy. But, since these parts are suspected to be saturated in at least one other image, they are omitted to make the registration possible. The resulting mask is used to define the exposure invariant image $\overline{E}_j(p)$:

$$\overline{E}_j(p) = E_j(p) \cdot M(p), \quad (5)$$

where $M(p)$ is the binary mask that is zero at the discarded pixel locations. Applying this transform ensures that the input images without the invalid pixels only differ in motion and not in intensity.

3.3. Rotation and Shift Estimation

For the registration of the exposure invariant images, we used the image registration algorithm proposed by Vandewalle et al.³ This method uses a frequency domain algorithm to estimate the motion parameters between the reference image and each other image. Only planar motion parallel to the image plane is allowed, as such motion can be well described in frequency domain. The motion can be described as a function of three parameters: horizontal and vertical shifts Δx_1 and Δx_2 and a planar rotation angle φ .

Let $\overline{\mathbf{E}}_1(u)$ and $\overline{\mathbf{E}}_2(u)$ be the Fourier transformation of two images $\overline{E}_1(p)$ and $\overline{E}_2(p)$, where u is the frequency index. The use of a frequency domain method allows us to estimate the shifts and rotations separately. A shift in spatial domain corresponds to a (linear) phase shift in frequency domain, while a rotation in spatial domain is a rotation of the entire Fourier transform (both phase and amplitude) over the same angle in frequency domain. We can therefore estimate the rotation first using the amplitudes of the Fourier transform, and then estimate the shifts using the phases after cancellation of the rotation.

To estimate the rotation angle φ between two images, we compute the frequency content $H_j(\alpha)$ as a function of the angle α by integrating the amplitude of the Fourier transform over radial lines:

$$H_j(\alpha) = \int_{\alpha - \Delta\alpha/2}^{\alpha + \Delta\alpha/2} \int_0^\infty |\overline{\mathbf{E}}_j(r, \theta)| d_r d\theta, \quad (6)$$

where $|\overline{E}_j(r, \theta)|$ is the Fourier transform image expressed in polar coordinates. The rotation angle φ can then be found as the value for which the correlation between $H_1(\alpha)$ and $H_2(\alpha)$ is maximized.

After inverting the rotation between the two images, the horizontal and vertical shifts can be estimated from the linear phase difference between the two resulting images. To perform a robust estimation, we compute a least squares fit of a plane through the phase difference.

In order to make our registration algorithm robust against aliasing in the input images, we only use the low frequency part of the spectrum for the registration. These frequencies are less aliased, and typically also less corrupted by noise. We perform such a pairwise registration for each of our exposure invariant images with respect to a reference image. This gives us a series of motion parameters, one for each image that we write $\Delta x_{1,j}$ and $\Delta x_{2,j}$ for the horizontal and vertical shift of image j and φ_j for the rotation. The reference image will have zero values.

3.4. Interpolation

Now that we have estimated the motion parameters, we can interpolate a high resolution image from the images E_j . First, we compute the pixel coordinates for each of the images on the high resolution grid using the alignment parameters. Next, we compute the pixel values on a high resolution, uniform grid of pixels from this nonuniform set of pixels using a Delaunay triangulation and bicubic interpolation method. For example, using four input images, we get

$$I_{HDR,SR}(p) = \psi(E_1(p), \Delta x_{1,1}, \Delta x_{2,1}, \varphi_1, E_2(p), \Delta x_{1,2}, \Delta x_{2,2}, \varphi_2, E_3(p), \Delta x_{1,3}, \Delta x_{2,3}, \varphi_3, E_4(p), \Delta x_{1,4}, \Delta x_{2,4}, \varphi_4), \quad (7)$$

where ψ is the interpolation function and $I_{HDR,SR}$ is the final high dynamic range and high resolution image. A tone mapping method should be applied on this image to have an appropriate image for viewing.

4. EXPERIMENTAL RESULTS

4.1. Data Gathering

Two different cameras were used for this project: a Sigma SD10 digital still camera and a Logitech QuickCam Pro webcam, and several scenes were captured. The exposure time was changed between acquisitions and a small camera motion was allowed. The motion between pictures was in the range of (-10, 10) pixels of horizontal and vertical shifts and (-1, 1) degrees of rotation.

4.2. Evaluation of the Camera Response Function Estimation

Since motion is allowed between acquisitions, our algorithm estimates the camera response function from the image histograms, which are invariant with motion. Here, we evaluate the benefit of this feature. First, we capture three images with different exposure times using a tripod so that there is no camera motion. Then, we create another set of images by shifting and rotating the first set of images using Matlab. Different shift and rotation values were tested as shown in Table 1 and several scenes were used for this experiment. As all the tested scenes led approximatively to the same results, we will only discuss one of them. We compute the CRF with Debevec et al.'s original algorithm⁸ and compare it to the CRF obtained with our histogram-based method. For each method, we compute the difference ϵ_{crf} between the CRF obtained with the set of images having no motion and the set of images where we artificially add motion. The error ϵ_{crf} is computed as follows:

$$\epsilon_{crf}(f_o, f_m) = \frac{1}{3} \sum_{c=1}^3 \frac{1}{256} \sum_{i=1}^{256} (f_{c,o}(i) - f_{c,m}(i))^2, \quad (8)$$

where c indicates the color channel, f_o is the estimated CRF for the original motionless image, and f_m is the estimated CRF for the rotated and shifted images.

Table 1 presents the result of this experiment for increasing motion parameters. The results show that the histogram-based method makes the CRF estimation more robust to shifts and rotations. Moreover, it shows that the larger the motion is, the larger the error becomes when the original method by Debevec et al. is applied. Figure 3 shows the estimated CRF curves for the red channel of the Sigma SD10 camera using the original and the proposed histogram-based method.

Table 1. Difference of CRF obtained when input images are shifted and rotated, with Debevec et al.’s method (original) and with our histogram-based version.

shift pixel	rot. angle degree	ϵ_{crf} with histogram	ϵ_{crf} original
$\pm[0.1, 0.1]$	± 0.1	0.0053	0.0094
$\pm[0.2, 0.2]$	± 0.2	0.0055	0.0120
$\pm[0.3, 0.3]$	± 0.3	0.0056	0.0143
$\pm[0.4, 0.4]$	± 0.4	0.0051	0.0183
$\pm[0.5, 0.5]$	± 0.5	0.0057	0.0238




Figure 3. Comparison of the CRF estimation (red color channel) for a set of motionless images and the same images artificially shifted and rotated ($[0.3 \ 0.3]$ pixel shift and 0.3 degree rotation). (a) Results for our histogram-based method. (b) Results for Debevec et al.’s method.

4.3. Exposure Invariant Transformation

Next, we evaluate our exposure invariant transform. To do so, we capture images of different exposure times using a tripod to avoid camera motion between the images. We then apply an artificial random motion on these images and compute the exposure invariant transform to obtain the exposure invariant masks M_j . Figure 4 shows two images taken with different exposure times and their corresponding exposure invariant masks. To




Figure 4. Two images taken with different exposure times and their corresponding exposure invariant masks. (a) First image (long exposure time). (b) Exposure invariant mask for (a). (c) Second image (short exposure time). (d) Exposure invariant mask for (c).

evaluate our transform and show that the masks M_j are exposure invariant, we take the masks obtained for two images with different exposure, invert the shift and rotation that was applied artificially, and count how many pixels are different. This experiment was repeated over several images. The results show that the exposure invariant masks of aligned images are 97% similar on average.

4.4. Rotation and Shift Estimation

Motion estimation is performed using the algorithm developed by Vandewalle et al.³ Here, we evaluate the performance of this algorithm when used on differently exposed images. Once again, we take a set of images with different exposure times using a tripod to prevent camera motion. Then, we artificially apply motion on these images using random motion parameters. We tested the images separately for rotation and shifts. Different rotations on different images resulted in an average error of 0.06 degrees. We proceeded similarly to estimate the shift error. We obtained average shift errors of 0.19 and 0.22 pixels for the vertical and horizontal shift estimation, respectively.

4.5. Final Result

Once the motion parameters are estimated, the high resolution, high dynamic range images are reconstructed using the algorithm described in Section 3.4. Figure 5 and Figure 6 show the results for two scenes. Figure 5 is obtained from a set of four input images that were taken with exposure times 1/4, 1/4, 1/100, and 1/100 seconds. The top figure shows the final result, the center figure is the low resolution, low dynamic range image captured with a long exposure time (1/4 s) and the bottom figure is the low resolution, low dynamic range image

taken with a short exposure time (1/100 s). The images on the right show a detail of the images on the left. The resulting image contains more details than any of the low dynamic range and low resolution images. This can be seen in the flower and the dress. In the detail images, we see that the spades shape is much better distinguishable in the high resolution version. Figure 6 shows an example of a high dynamic range scene. The input images were taken with exposure times 1/200 (right), and 1/800 (center) seconds. In addition to the increase in resolution we can see that the saturated part that contain no details in the center image does contain data in the result image (left).

Figure 7 shows another example of increased resolution. The top image shows one of the low resolution images. In the low resolution image, aliasing is visible in the high frequency parts of the image. Our method increases the resolution and thus decreases the aliasing in the resulting image (Figure 7(b)).

This algorithm can also be applied to a movie taken with frames of alternating exposure times, such that a high resolution and high dynamic range movie can be generated using a standard video camera. To test the efficiency of our algorithm for this purpose, we created movies composed from dark and bright frames by moving a window over our captured scene. The even frames are taken from the low exposure image and the odd ones are taken from the short exposure image. Between both images, small shifts and rotations were made. We applied this procedure in some simulations, and obtained high resolution and high dynamic range movies. The results are shown online at http://lcavwww.epfl.ch/reproducible_research/.

5. CONCLUSION

We presented a method to construct a high resolution and high dynamic range image from a set of low resolution and limited dynamic range images. We used the algorithm designed by Vandewalle et al.³ for super-resolution and an adapted version of the method proposed by Debevec et al.⁸ for estimating the camera response function. The results show a clear improvement in resolution and dynamic range. We also showed the improvements in the CRF estimation in comparison with Debevec et al.'s original method.⁸ We can also apply our method to use differently exposed input images in a super-resolution approach, or to acquire high dynamic range images using a handheld camera.

REFERENCES

1. S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," *IEEE Signal Processing Magazine* **20**, pp. 21–36, May 2003.
2. S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Advances and challenges in super-resolution," *International Journal of Imaging Systems and Technology* **14**, pp. 47–57, Aug. 2004.
3. P. Vandewalle, S. Süsstrunk, and M. Vetterli, "A Frequency Domain Approach to Registration of Aliased Images with Application to Super-Resolution," *EURASIP Journal on Applied Signal Processing, Special Issue on Super-Resolution Imaging* **2006**, 2006. Article ID 71459, 14 pages.
4. B. Marcel, M. Briot, and R. Murrieta, "Calcul de translation et rotation par la transformation de Fourier," *Traitement du Signal* **14**(2), pp. 135–149, 1997.
5. L. Lucchese and G. M. Cortelazzo, "A noise-robust frequency domain technique for estimating planar rotations," *IEEE Transactions on Signal Processing* **48**, pp. 1769–1786, June 2000.
6. D. Keren, S. Peleg, and R. Brada, "Image sequence enhancement using sub-pixel displacement," in *Proceedings IEEE conference on computer vision and pattern recognition*, pp. 742–746, June 1988.
7. D. Robinson, S. Farsiu, and P. Milanfar, "Optimal Registration of Aliased Images Using Variable Projection with Applications to Superresolution," *submitted to The Computer Journal*, 2006.
8. P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *Proc. ACM SIGGRAPH 97, Annual Conference on Computer Graphics*, pp. 369–378, (Los Angeles, CA), August 1997.




(a)

(b)



(c)

(d)



(e)

(f)

Figure 5. Result using our approach. (a) Output result. (c) Brightest input image with $\frac{1}{4}$ s exposure time. (e) Darkest input image with $\frac{1}{100}$ s exposure time. Details of the images in (a), (c), and (e) are shown in (b), (d), and (f), respectively.




Figure 6. Result using our approach. (a) Output result. (b) Brightest input image with $\frac{1}{400}$ s exposure time. (c) Darkest input image with $\frac{1}{1600}$ s exposure time.

9. S. Mann and R. W. Picard, "On being 'undigital' with digital cameras: extending dynamic range by combining differently exposed pictures," in *Proc. IS&T's 48th Annual Conference*, pp. 422–428, (Cambridge, MA), May 1995.
10. S. Mann, "Comparametric equations with practical applications in quantigraphic image processing," *IEEE Transactions on Image Processing* **9**, pp. 1389–1406, Aug. 2000.
11. M. A. Robertson, S. Borman, and R. L. Stevenson, "Dynamic Range Improvement Through Multiple Exposures," in *Proceedings IEEE International Conference on Image Processing*, **3**, pp. 159–163, Oct. 1999.
12. D.-W. Kim and K.-S. Hong, "High Dynamic Range Global Mosaic," *Lecture Notes in Computer Science (Proc. Asian Conference on Computer Vision)* **3851**, pp. 744–753, Jan. 2006.
13. Y. Y. Schechner and S. K. Nayar, "Generalized Mosaicing: High Dynamic Range in a Wide Field of View," *International Journal of Computer Vision* **53**(3), pp. 245–267, 2003.
14. B. K. Gunturk, "Handling Exposure Time in Multi-Frame Image Restoration," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, **2**, pp. 865–868, May 2005.
15. A. A. Goshtasby, "Fusion of multi-exposure images," *Image and Vision Computing* **23**, pp. 611–618, 2005.
16. S. G. Narasimhan and S. K. Nayar, "Enhancing resolution along multiple imaging dimensions using assorted pixels," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27**, pp. 518–530, April 2005.
17. B. K. Gunturk and M. Gevrekci, "High-resolution image reconstruction from multiple differently exposed images," *IEEE Signal Processing Letters* **13**, pp. 197–200, Apr. 2006.
18. G. Ward, "Fast, robust image registration for compositing high dynamic range photographs from hand-held exposures," *Journal of Graphics Tools* **8**(2), pp. 17–30, 2003.



(a)



(b)

Figure 7. Aliasing is reduced using our approach. (a) Brightest input image. (b) Output image of our algorithm.