Double resolution from a set of aliased images

Patrick Vandewalle[1], Sabine Süsstrunk[1] and Martin Vetterli[1,2]

[1]LCAV - School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne (EPFL)
CH-1015 Switzerland
[2]Department of Electrical Engineering and Computer Science
UC Berkeley
Berkeley CA 94720, USA

ABSTRACT

In this paper, we present a super-resolution method to approximately double image resolution in both dimensions from a set of four low resolution, aliased images. The camera is shifted and rotated by small amounts between the different image captures. Only the low frequency, aliasing-free part of the images is used to find the shift and rotation parameters. When the images are registered, it is possible to reconstruct a higher resolution, aliasing-free image from the four low resolution images using cubic interpolation. We applied our algorithm in a simulation, where all parameters are known and controlled, as well as in a practical experiment using images taken with a real digital camera. The results obtained in both tests prove the validity of our method.

Keywords: super-resolution imaging, subpixel planar motion estimation, shift and rotation estimation, aliased images

1. INTRODUCTION

For many applications, digital cameras are replacing their analog counterparts. Recently, prices of digital imaging equipment have decreased sharply. At the same time, the image quality of digital cameras is continuously increasing, showing for most applications no more perceivable difference between images taken with analog and digital cameras. It is therefore no surprise that digital cameras can now be found in many applications, in all forms, colors and prices. Meanwhile, the transition from analog to digital devices introduces new problems to be dealt with, such as sensor noise, aliasing, demosaicing, etc.

One of the consequences of the use of digital cameras is aliasing, which can severely degrade the quality of a digital image. These artifacts do not occur in images taken with traditional film cameras, because they are caused by the sampling process in digital cameras. If the sampling frequency is too low compared to the maximal signal frequency, aliasing effects appear. Then too little information about the sampled signal (in our case a 2D image) is available in order to be able to reconstruct the original signal unambiguously.

Intuitively, the solution to this problem is to increase the sampling frequency, such that the Nyquist criterium is satisfied and, consequently, aliasing is avoided. But this solution doesn't work for digital cameras, because we cannot change their sampling frequency. It is fixed with the choice of a CCD sensor and cannot be modified depending on the scene. The solution that is now mostly chosen by digital camera manufacturers is to use an optical system that cuts off most frequencies above half the sampling frequency (and thus satisfies the Nyquist criterium again). While this certainly eliminates all aliasing artifacts, it also removes part of the image signal. Another solution could be to add more samples from a second sampling set (image) of the same signal (scene) at different locations.

This is the super-resolution approach: a set of images of the same scene is combined to form a new, more densely sampled image of the scene. This idea was first introduced by Tsai and Huang[1] in 1984. A good overview of the super-resolution literature is given by Borman and Stevenson[2] and Park et al.[3]

Further author information: (Send correspondence to Patrick Vandewalle)
E-mail: {Patrick.Vandewalle, Sabine.Susstrunk, Martin.Vetterli}@epfl.ch, Telephone: +41 21 6935125, Fax: +41 21 6934312

Most super-resolution approaches can be subdivided in two (main) parts: image registration and image reconstruction. Image registration is performed first, in order to align the different images as precisely as possible. A sub-pixel precision is required for the alignment, because it is exactly this in-between pixel information that will be used to increase image resolution and remove aliasing.

Different approaches exist for estimating motion between two images. Frequency domain methods for estimating shifts in the image plane are based on the fact that a shift in spatial domain can be expressed as a phase shift in frequency domain. Two shifted images differ only by a linear phase difference.[4,5] These methods can be extended to include (planar) rotation and scale using polar coordinates,[6] with the advantage that shift, rotation and scale can be estimated separately. The main limitation of frequency domain methods is that they are restricted to global shifts and rotations in the image plane, and scale. Our algorithm uses a frequency domain method, with a rotation estimation method that does not need a conversion to polar coordinates. Spatial domain registration methods allow the use of more general motion models, like homographies. Irani et al.[7] describe a spatial domain method that allows for different motions in the same image. Capel and Zisserman[8] use a maximum likelihood method to determine the homography parameters from corresponding feature vectors.

In a second step, a high resolution image is reconstructed from the set of registered low resolution images. This can be done using nonuniform interpolation. Another possibility is to apply a method based on projection onto convex sets (POCS), where the final image is assumed to belong to different convex sets with a (possibly non-unique) intersection. Keren et al.[9] describe an iterative method starting from an initial guess of the high resolution image. This image is then improved by simulating the imaging process and minimizing the difference between the simulated and observed low resolution images. In our approach, we use cubic interpolation for reconstruction of the high resolution image.

Some methods do not compute the registration parameters explicitly. The optimal high resolution image is directly estimated using a (constrained) least squares or a maximum a posteriori (MAP) approach.[10,11] A regularization parameter is often used in these algorithms to avoid multiple solutions.

This paper is structured as follows. In Section 2 we describe our registration method. First, general concepts on the separation of rotation and shift estimation are developed. Then, we discuss the rotation and the shift estimation algorithms, respectively. And finally, some modifications for registration of aliased images are proposed. In Section 3, the reconstruction method is described, and in Section 4, the complete algorithm is presented. We show results on simulated and aliased images captured with a digital camera in Section 5. In Section 6, we discuss possible applications of our method, and Section 7 concludes the article.

2. IMAGE REGISTRATION

We use a frequency domain method to register the images. Therefore, we assume that the motion between two images will be a planar motion in a plane parallel to the image plane. Motion can then be described by three parameters: horizontal and vertical shift Δx and Δy and a rotation angle ϕ.

Using these parameters, we can describe a shifted and rotated version $f_2(\mathbf{x})$ of a signal $f_1(\mathbf{x})$ as

$$f_2(\mathbf{x}) = f_1(\mathbf{R}(\mathbf{x} + \Delta\mathbf{x})), \quad \text{with } \mathbf{x} = \begin{bmatrix} x \\ y \end{bmatrix}, \Delta\mathbf{x} = \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}, \mathbf{R} = \begin{bmatrix} \cos\phi & -\sin\phi \\ \sin\phi & \cos\phi \end{bmatrix}. \quad (1)$$

Note that the image has been first shifted and then rotated. In order to find the same values of Δx, Δy and ϕ, we will need to cancel first rotation and then shift. We can describe this in the Fourier domain as

$$\begin{aligned} F_2(\mathbf{u}) &= \iint_{\mathbf{x}} f_2(\mathbf{x}) e^{-j2\pi \mathbf{u}^T \mathbf{x}} d\mathbf{x} \\ &= \iint_{\mathbf{x}} f_1(\mathbf{R}(\mathbf{x} + \Delta\mathbf{x})) e^{-j2\pi \mathbf{u}^T \mathbf{x}} d\mathbf{x} \\ &= e^{j2\pi \mathbf{u}^T \Delta\mathbf{x}} \iint_{\mathbf{x}'} f_1(\mathbf{R}\mathbf{x}') e^{-j2\pi \mathbf{u}^T \mathbf{x}'} d\mathbf{x}', \end{aligned} \quad (2)$$

with $\mathbf{x}' = \mathbf{x} + \Delta\mathbf{x}$, and $F_2(\mathbf{u})$ the Fourier transform of $f_2(\mathbf{x})$. Using another transformation $\mathbf{x}'' = \mathbf{R}\mathbf{x}'$, we can therefore compute

$$
\begin{aligned}
|F_2(\mathbf{u})| &= \left| e^{j2\pi\mathbf{u}^T\Delta\mathbf{x}} \iint_{\mathbf{x}'} f_1(\mathbf{R}\mathbf{x}')e^{-j2\pi\mathbf{u}^T\mathbf{x}'}d\mathbf{x}' \right| \\
&= \left| \iint_{\mathbf{x}'} f_1(\mathbf{R}\mathbf{x}')e^{-j2\pi\mathbf{u}^T\mathbf{x}'}d\mathbf{x}' \right| \\
&= \left| \iint_{\mathbf{x}''} f_1(\mathbf{x}'')e^{-j2\pi\mathbf{u}^T(\mathbf{R}^T\mathbf{x}'')}d\mathbf{x}'' \right| \\
&= \left| \iint_{\mathbf{x}''} f_1(\mathbf{x}'')e^{-j2\pi(\mathbf{R}\mathbf{u})^T\mathbf{x}''}d\mathbf{x}'' \right|,
\end{aligned}
\tag{3}
$$

which does not depend on the shifts and is a rotated version of $|F_1(\mathbf{u})|$. It is therefore possible to estimate the relative rotation angle between two images first. We can then apply the inverse rotation to the Fourier transform. The (phase) shift can be computed from the resulting image and the reference image.

2.1. Rotation estimation

To estimate the rotation angle between two images, we can compute the correlation between the reference image and rotated versions of the second image. The estimated rotation angle is the angle for which a maximum correlation is obtained. The main disadvantages of this method are its computational cost and the interpolation errors that result from every rotation.

We developed a different method where only a one-dimensional correlation is needed. The average signal energy is computed for every angle of the Fourier transform of the two images. This gives us the energy in the two images as a function of the angle. We can then find the rotation angle as the maximum of the circular correlation between the two functions. This is a much faster method as it does not require as many correlations or interpolations.

The lowest frequency values are discarded for the computation of the average energy, because they are generally very large and very coarsely sampled in terms of angle. This would bias the results. Only a disc region with radius from 0.1ρ to ρ, with ρ the image radius (half of the image size), is used in order to have similar numbers of values at the different angles (see Figure 1).

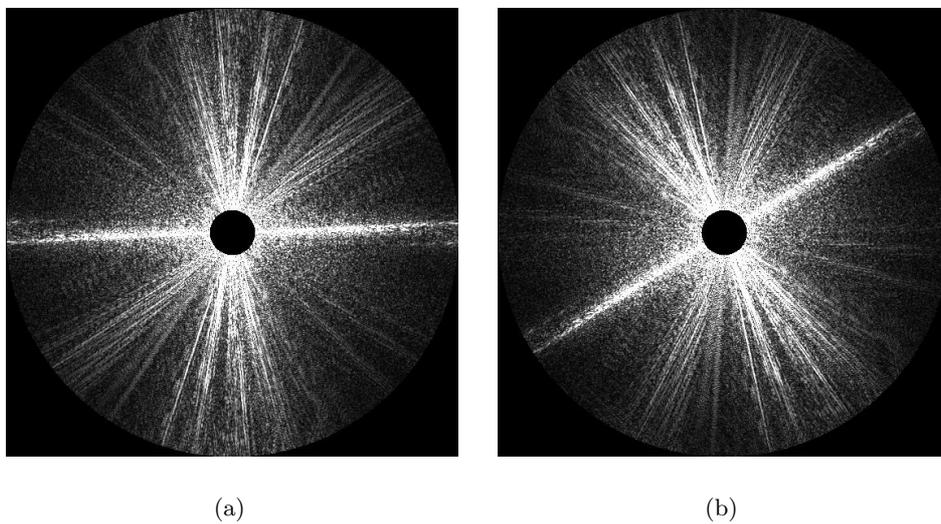(a)   (b)

Figure 1. Fourier transform of an original image (a) and after rotation over $30°$ (b). Only a disc with radius $0.1\rho < r < \rho$ (with r the radius and ρ half of the image size) is considered for the rotation estimation.

2.2. Shift estimation

A shift in the image plane corresponds to a phase shift in frequency domain. For an image $f_1(\mathbf{x})$ and its shifted version $f_2(\mathbf{x})$, with Fourier transforms $F_1(\mathbf{u})$ and $F_2(\mathbf{u})$, respectively, we can express this as

$$F_2(\mathbf{u}) = e^{j2\pi\mathbf{u}^T\Delta\mathbf{x}} \iint_{\mathbf{x}} f_1(\mathbf{x})e^{-j2\pi\mathbf{u}^T\mathbf{x}}d\mathbf{x}$$
$$= e^{j2\pi\mathbf{u}^T\Delta\mathbf{x}}F_1(\mathbf{u}). \tag{4}$$

Thus, the difference between the phase of the two Fourier transforms is a plane in u-v space. Horizontal and vertical shifts can therefore be directly computed as the slope of this plane in horizontal and vertical direction, respectively. Because of noise in the images, differences because of non-overlapping parts, etc. this phase difference is only approximately a plane. We solve this problem by estimating the plane parameters from the data with a least squares method to make the estimates more robust to errors.

2.3. Aliasing and circular extension

The registration methods described above work for perfect, noise-free images. In reality, however, images are degraded by noise and aliasing. This often causes failure of the image registration.

Assume the images, with maximal frequency $\mathbf{f}_{max}$, are sampled at frequency $\mathbf{f}_s$

$$\text{with } \mathbf{f}_{max} < \mathbf{f}_s < 2\mathbf{f}_{max} \text{ and all } \mathbf{f} = \left[\begin{array}{c} f_x \\ f_y \end{array} \right].$$

This setup does not satisfy the Nyquist criterium ($\mathbf{f}_s > 2\mathbf{f}_{max}$), and therefore aliasing artifacts will be present in the sampled images. But, as $\mathbf{f}_s > \mathbf{f}_{max}$, the signal is free of aliasing and thus the same in all images (up to shift and rotation effects) at frequencies

$$-\mathbf{f}_s + \mathbf{f}_{max} < \mathbf{f} < \mathbf{f}_s - \mathbf{f}_{max}.$$

We can therefore do the image registration based on these frequencies and use the complete frequency content to reconstruct the high resolution image accurately. This principle has already been applied to image registration by Kim and Su[4] and to super-resolution by Vandewalle et al.[12]

Another cause of artifacts in the Fourier transform of images is the fact that they are (in general) not circular. When computing the FFT of an image, it is circularly extended, which causes strong artificial edges in the spatial domain image. In the frequency domain, this mainly causes artifacts around the two axes. In order to eliminate this effect, we multiply the images by a (Tukey) window prior to registration. The border lines of the image all have zero value and the image is successfully circularized.

3. IMAGE RECONSTRUCTION

The high resolution image is reconstructed using a cubic interpolation method. First, the coordinates of the pixels are computed for all images in the coordinate frame of the first image. Next, a regular, higher resolution grid is constructed for the new image. And finally, the image values on this grid are interpolated from the known (irregularly sampled) pixel values using cubic interpolation.

How many images are required to increase the resolution by a certain factor? Intuitively, it seems that with more images, we can always construct a better, higher resolution image. However, there are some limitations. Because we assumed the sampling frequency to be $\mathbf{f}_s > \mathbf{f}_{max}$, the actual resolution can be increased by a factor of two in each dimension. It cannot be increased more, because higher frequencies are not present in the measured images. When the image resolution is increased by a factor two in both dimensions, this results in an image with four times the original number of pixels. In frequency domain, we could say that in order to compute all frequency components of the new image, four times as many measurements would be needed. This is not generally true for reconstruction using cubic interpolation, but we assume four images are still a good trade-off between high precision (many images) and low complexity (few images).

4. ALGORITHM

The different methods described in earlier sections are combined here in the description of the complete super-resolution algorithm.

1. All four images i_1, i_2, i_3 and i_4 are multiplied with a Tukey window to minimize border effects. The resulting images are called $i_{1,w}, i_{2,w}, i_{3,w}$ and $i_{4,w}$.

2. Estimate rotation

   (a) Compute the Fourier transforms $I_{1,w}, I_{2,w}, I_{3,w}$ and $I_{4,w}$ of the images $i_{1,w}, i_{2,w}, i_{3,w}$ and $i_{4,w}$.

   (b) Compute the average energy at all different angles, up to a precision of 0.1 (degrees). Do this for $I_{1,w}, I_{2,w}, I_{3,w}$ and $I_{4,w}$.

   (c) Correlate the energy functions of $I_{2,w}, I_{3,w}$ and $I_{4,w}$ with the one from $I_{1,w}$.

   (d) The estimated rotation angle $\hat{\phi}$ is the maximum of this correlation function.

3. Rotate the images $i_{2,w}, i_{3,w}$ and $i_{4,w}$ over $-\hat{\phi}$ and use these rotated images $i_{1,w}, i_{2,wr}, i_{3,wr}$ and $i_{4,wr}$ for the shift estimation.

4. Estimate horizontal and vertical shift

   (a) Compute the Fourier transforms $I_{1,w}, I_{2,wr}, I_{3,wr}$ and $I_{4,wr}$ of the images $i_{1,w}, i_{2,wr}, i_{3,wr}$ and $i_{4,wr}$.

   (b) Compute the phase difference between the images as $\angle(I_{2,wr}/I_{1,w})$, $\angle(I_{3,wr}/I_{1,w})$ and $\angle(I_{4,wr}/I_{1,w})$.

   (c) Find the least squares solution to fit a plane through the phase values for the lowest 5% of frequencies.

   (d) Derive the horizontal and vertical shifts from the plane parameters.

5. Compute the coordinates of the images i_2, i_3 and i_4 in the coordinate frame of i_1 using the estimated rotation and shift parameters.

6. Interpolate the values on a regular $2Mx2N$ grid using cubic interpolation from the set of image points in i_1, i_2, i_3 and i_4 (where MxN is the size of the original images).

5. RESULTS

The method was tested in a simulation setup as well as in a practical experiment. The goal of the simulation was to evaluate the performance in a controlled environment with exact knowledge about the original high resolution image and about the shift and rotation values. This enabled us to evaluate the registration and reconstruction errors. In the practical experiment, the algorithm is tested on aliased pictures taken with a real digital camera.

5.1. Simulation

For the simulation, we started from an original high resolution (1704x1704 pixels) image, which was taken as equivalent for continuous space (Figure 3(a)). In order to have exactly the same content in all low resolution images, the original image was multiplied by a Tukey window and zero-padded. From this original image, four low resolution images (221x221 pixels) were constructed by shifting and rotating the original and then downsampling by a factor eight (Figure 2 and Figure 3(b)). Prior to downsampling, the images are low-pass filtered such that 5% of the frequency domain of the resulting low resolution images is free of aliasing in horizontal and vertical directions.

These four low resolution images are then used as input to the image registration algorithm. When only shifts are applied, the image registration is perfect. When combined shifts and rotations are applied, the rotation angle is accurately estimated, and the shifts are well approximated. The errors in the shift estimates are due to the interpolation approximations made when rotating the images.

Finally, the four images and the registration information are used to reconstruct a high resolution image (Figure 3(e) and 3(f)). The reconstructed image has a mean squared error of 2.5e-3 , where all pixel values are scaled between 0 and 1 (Figure 3(g)). The precision of the image registration was sufficient to have a good reconstruction (see Table 1).
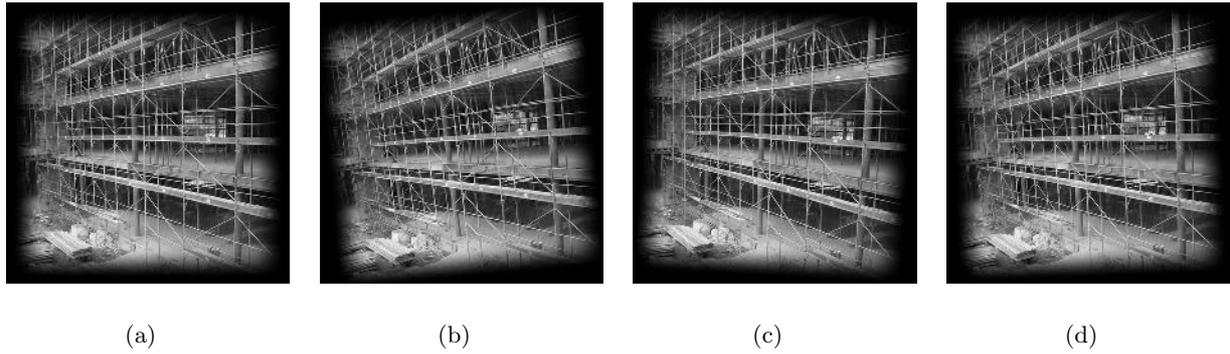
| | (a) | (b) | (c) | (d) |

Figure 2. Four low resolution images to be registered. Actual parameter values Δx, Δy and ϕ and estimated values Δx_{est}, Δy_{est} and ϕ_{est} are given for all four images in Table 1.

	Δx	Δy	ϕ	Δx_{est}	Δy_{est}	ϕ_{est}
Image 1 (ref)	0	0	0	0	0	0
Image 2	0.625	0.25	5	0.64	0.17	5
Image 3	0.25	0.75	-3	0.19	0.68	-3
Image 4	0.375	0.125	2	0.35	0.14	2

Table 1. Comparison between the exact motion parameters (Δx, Δy and ϕ) and the estimated parameters (Δx_{est}, Δy_{est} and ϕ_{est}) in the simulation.

5.2. Experiment

The scene we used in the practical experiment was a flat glass plate with black print, positioned parallel to the camera image plane. The plate was illuminated from the back. A black and white digital camera was used to capture the scene. The camera was mounted on a solid tripod, allowing for separate horizontal and vertical shifts and rotations in the image plane. The optical system for the camera was chosen such that the images have aliasing artifacts.

A set of four images was taken using this setup (Figure 4(a)). The camera was shifted in both directions and rotated by a small amount between the different images. We estimate that 60% of the frequency domain of the images is free of aliasing. The image registration algorithm was then directly applied to this set of images, registering them all to the first image. The estimated registration parameters are

$$\text{Image 2 - Image 1:} \quad \Delta x_{est} = 9.24, \quad \Delta y_{est} = -3.84, \quad \phi = 0.9°$$
$$\text{Image 3 - Image 1:} \quad \Delta x_{est} = 9.89, \quad \Delta y_{est} = -2.29, \quad \phi = 1.1°$$
$$\text{Image 4 - Image 1:} \quad \Delta x_{est} = 10.41, \quad \Delta y_{est} = -5.04, \quad \phi = 1.1°.$$

From the four registered images, a high resolution image was reconstructed (Figure 4(b)). As can be seen in Figure 4(b), the images have been precisely aligned for the construction of the high resolution image and most aliasing has been accurately removed. Remark that in the printing process of this article, the images are probably resampled. Depending on the sampling rate, this may have reintroduced some aliasing artifacts.

6. DISCUSSION AND FUTURE WORK

High accuracy in image registration is crucial in super-resolution imaging. If good motion estimates at a subpixel level are not available, image reconstruction will almost certainly result in a blurred, low quality image. It is then better to interpolate a high resolution image from one single low resolution image, rather than incorrectly using additional information from other images. Accurate subpixel shift estimates and sub-degree rotation estimates as in the registration methods described above are necessary.

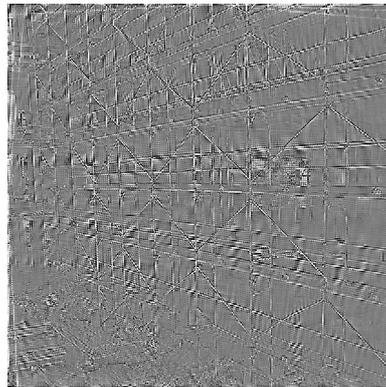**Figure 3.** (a) Original high resolution image. (b) Low resolution image used as input for the super-resolution algorithm. (c) High resolution reconstruction target. (d) Detail of the reconstruction target. (e) Reconstructed high resolution image (MSE=2.5e−3). (f) Detail of the reconstructed image. (g) Difference between the reconstruction target and the reconstructed image (neutral gray corresponds to zero error, the values have been amplified for better visibility).
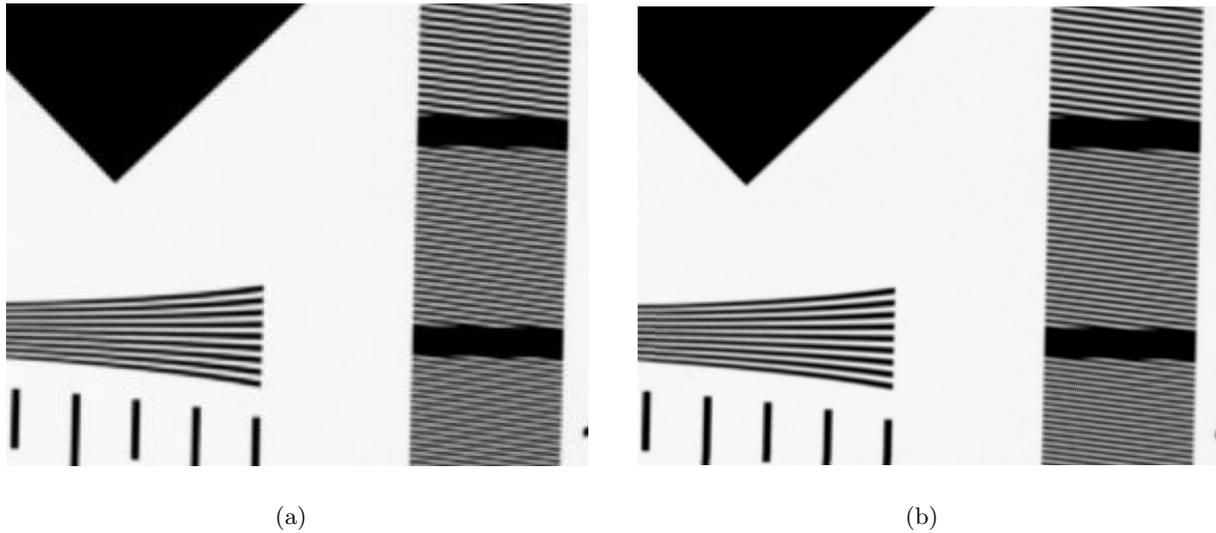
(a)　　　　　　　　　　　　　　　　　　　(b)

Figure 4. (a) One of the images used as input for the super-resolution algorithm. (b) Reconstructed high resolution image, from which aliasing has been accurately removed.

This technique could be applied, for example, in consumer digital cameras. The size of the CCD still largely determines the price of digital cameras. Thus, it would be very interesting to use a digital camera with a smaller CCD and create higher resolution images with it, by taking multiple shots. Assume the user holds his digital camera manually while (manually or automatically) taking a series of four images. There will be small shifts and rotations between the images, because the camera is only held more or less in the same position. These small motions between the four images could now be first estimated and then be used to construct an image with double resolution in horizontal and vertical direction. Because of the small amount of motion between the different images, the non-overlapping part between the images is also very small. The frequency domain signals will also be very similar, and consequently the motion estimation will be more precise.

The optical systems that are mostly used nowadays limit the application of this super-resolution technique in digital cameras. Because aliasing artifacts are very disturbing, the optical system of most digital cameras is designed such that all high frequencies are removed. The entire image is slightly blurred, but no more aliasing is visible. This of course causes our algorithm to fail, because it relies exactly on this high frequency information. In our method, we assumed for both image registration and reconstruction that the image was sampled using a Dirac sampling kernel. This is, of course, very different from the Gaussian sampling kernel, which is often referred to as a good model for the sampling process in digital cameras. It would therefore be very interesting for future research to include this more realistic sampling model in our super-resolution method.

It would also be very interesting to study the possibilities of this method for demosaicing applications. In most (consumer) color digital cameras, one CCD is used with a color filter array (CFA) to measure red, green and blue channels of the image. With the Bayer pattern, which is often used for the CFA, there are twice as many green samples (pixels) as red or blue samples. A demosaicing algorithm is applied to this sampled image to create an RGB image with all three color values at every pixel location. As the resolution of the red and blue image is upsampled by two in this process, (color) aliasing can occur. If we capture multiple shots of the same scene, the super-resolution algorithm can be applied to those red and blue images in order to remove all aliasing from the resulting RGB image.

7. CONCLUSIONS

In this paper, we developed a new super-resolution method to almost double the resolution of images in both dimensions. We started from a set of four aliased, low resolution images with small shifts and rotations between

them. Shift and rotation were estimated from the low-pass, aliasing-free part with an efficient frequency domain method. Using this registration information, we were able to reconstruct a high resolution image by applying a cubic interpolation method.

Our algorithm was tested in a simulation where all the parameters are controlled, and in a practical experiment. Both experiments showed very good results. Aliasing has been accurately removed from the resulting images.

ACKNOWLEDGMENTS

REFERENCES

1. T. S. Huang, ed., *Advances in Computer Vision and Image Processing*, vol. 1, ch. 7, pp. 317–339. JAI Press, 1984.
2. S. Borman and R. Stevenson, "Spatial resolution enhancement of low-resolution image sequences - a comprehensive review with directions for future research," tech. rep., University of Notre Dame, 1998.
3. 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.
4. S. P. Kim and W.-Y. Su, "Subpixel accuracy image registration by spectrum cancellation," in *Proceedings IEEE International Conference on Acoustics, Speech and Signal Processing*, **5**, pp. 153–156, April 1993.
5. H. S. Stone, M. T. Orchard, E.-C. Chang, and S. A. Martucci, "A fast direct Fourier-based algorithm for subpixel registration of images," *IEEE Transactions on Geoscience and Remote Sensing* **39**, pp. 2235–2243, October 2001.
6. 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.
7. M. Irani, B. Rousso, and S. Peleg, "Computing occluding and transparent motions," *International Journal of Computer Vision* **12**, pp. 5–16, February 1994.
8. D. Capel and A. Zisserman, "Computer vision applied to super resolution," *IEEE Signal Processing Magazine* **20**, pp. 75–86, May 2003.
9. 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.
10. S. Baker and T. Kanade, "Limits on super-resolution and how to break them," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24**, pp. 1167–1183, September 2002.
11. M. Elad and A. Feuer, "Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images," *IEEE Transactions on Image Processing* **6**, pp. 1646–1658, December 1997.
12. P. Vandewalle, S. Süsstrunk, and M. Vetterli, "Superresolution images reconstructed from aliased images," in *SPIE/IS&T Visual Communication and Image Processing Conference*, T. Ebrahimi and T. Sikora, eds., **5150**, pp. 1398–1405, July 2003.