

Registration of aliased images for super-resolution imaging

Patrick Vandewalle^a, Luciano Sbaiz^a, Sabine Süsstrunk^a and Martin Vetterli^{a,b}

^aSchool of Computer and Communication Sciences,
Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland;

^bDepartment of Electrical Engineering and Computer Science,
UC Berkeley, Berkeley CA94720, USA

ABSTRACT

Super-resolution imaging techniques reconstruct a high resolution image from a set of low resolution images that are taken from almost the same point of view. The problem can be subdivided into two main parts: an image registration part where the different input images are aligned with each other, and a reconstruction part, where the high resolution image is reconstructed from the aligned images.

In this paper, we mainly consider the first step: image registration. We present three frequency domain methods to accurately align a set of undersampled images. First, we describe a registration method for images that have an aliasing-free part in their spectrum. The images are then registered using that aliasing-free part. Next, we present two subspace methods to register completely aliased images. Arbitrary undersampling factors are possible with these methods, but they have an increased computational complexity. In all three methods, we only consider planar shifts. We also show the results of these three algorithms in simulations and practical experiments.

Keywords: super-resolution imaging, image registration, aliasing, undersampling

1. INTRODUCTION

The resolution, or resolving power, of a digital imaging device is mainly limited by the number of pixels on the sensor and the optical system. The goal of super-resolution imaging is to increase the resolution of an image by combining the information from multiple images taken from almost the same point of view. In this paper, we will study the super-resolution problem when the number of pixels is the limiting factor. In other words, the low resolution images are not lowpass filtered prior to sampling, and are therefore aliased.

Most super-resolution algorithms can be decomposed into two parts: an image registration part followed by a reconstruction. Very high accuracy is required in the registration (up to subpixel level) to be able to reconstruct the high resolution image correctly. Next, a robust reconstruction method is needed to undo the blur caused by the optical system and build a high resolution image from the set of irregularly spaced samples (pixels). We will mainly study the first problem, as an accurate registration is an essential requirement for a high precision reconstruction afterwards.

Such techniques can be applied in many domains. In digital photography, a user can hold his camera manually in approximately the same position while taking a series of pictures. The small movements between the different images allow to create a higher resolution image. A lower resolution, and thus cheaper digital camera can then be used to generate high resolution images. Similarly, in satellite imaging, small movements between different images can allow us to increase the precision of the remote sensing devices. In video surveillance, a super-resolution image created from a sequence of images can give much higher precision than individual images from the original sequence.

Remark that we use the term ‘resolution’ in the sense of resolving power: the power to discriminate details in an image. An image with ten times more pixels than another image does not necessarily have a higher resolving power. The super-resolution methods that are discussed here attempt to increase the resolving power, as opposed to for example a bilinear interpolation from a single image, which only adds pixels but does not improve resolving power.

The rest of this paper is structured as follows. The state of the art in super-resolution imaging, with an emphasis on image registration methods, is discussed in Section 2. Some new methods for registration of undersampled images are presented in Section 3, and the results of these methods in simulations and practical experiments are shown in Section 4. Advantages and disadvantages of the presented methods are discussed in Section 5, and finally, some conclusions are drawn in Section 6.

2. STATE OF THE ART

The super-resolution literature dates back to the algorithm described by Tsai and Huang¹ in 1984. They describe a frequency domain method that is applied to satellite imaging. The registration parameters are computed by searching the parameter values for which the reconstructed signal is bandlimited to a certain frequency. It is shown that for (small) deviations from these optimal values, the signal always has a larger bandwidth. The Fourier coefficients of the high resolution image can then be computed efficiently using an FFT-based algorithm.

A good overview of existing super-resolution methods is given by Borman and Stevenson.² More recently, special issues on super-resolution imaging appeared in *IEEE Signal Processing Magazine*³ and *EURASIP Journal on Applied Signal Processing*.⁴ Matlab implementations with a graphical user interface for certain super-resolution algorithms were also recently released by Farsiu et al.⁵ and Vandewalle et al.⁶

Image registration, which is the first part of most super-resolution methods, can be performed either in spatial or in frequency domain. A survey of image registration methods is given by Brown⁷ and by Zitova and Flusser.⁸ We summarize some of the most frequently used spatial domain techniques here, before we proceed to a more detailed overview of the frequency domain methods. One popular technique is to compute salient features and find a mapping between corresponding features in different images.^{9,10} Another approach uses a Taylor series approximation of the images. The motion parameters are the unknowns in the approximation, and they can be computed from the set of equations that can be derived from this approximation. Because Taylor series only give a good approximation for small offsets, these registration methods are generally applied iteratively using a Gaussian pyramid.^{11,12} Other methods compute an optical flow field,¹³ and can discriminate different motion patterns in a single image.

The frequency domain methods are limited to global motion in a plane parallel to the image plane. Planar shifts can be described with linear phase shifts in frequency domain, while rotations (in the image plane) in spatial domain correspond to rotations in frequency domain. Other, more general motion models are difficult to describe in frequency domain. However, frequency domain registration methods can be computationally very efficient, and offer a good framework to model aliasing.

If the input images are not undersampled, and their motion can be approximated by a planar shift, this shift can be computed as the linear phase difference between the two images.¹⁴ Planar rotation can be added, and is represented in frequency domain by a rotation over the same angle. The shift and rotation parameters can be estimated separately, because shift only affects the phase information, while rotation affects both phase and amplitude of the Fourier transform. Reddy and Chatterji¹⁵ and Marcel et al.¹⁶ applied this in their image registration algorithm. To estimate the rotation, they transform the Fourier domain image into polar coordinates, such that the rotation angle is transformed into a shift. Another rotation estimation method is presented by Lucchese and Cortelazzo.¹⁷ They compute the rotation angle by searching the zero crossings of the difference between the spectrum of one image and a mirrored version of the other image spectrum. The shift is estimated in the same way as Marcel et al.

If the low resolution images are aliased, these simple relations do not hold anymore. If part of the spectrum (generally the low frequencies) is free of aliasing, the shift parameters can still be estimated using only this part.¹⁸⁻²⁰ A rotation estimation method that is based on the low frequencies was added by Vandewalle et al.²⁰ This method is described in more detail in Section 3.2. Such an approach also gives a good approximation if the ratio between the signal power in the base spectrum and the power in the aliased spectrum copies is high. In that case, the aliasing can be considered as (part of the) noise.

If the aliasing cannot be neglected in any part of the spectrum, other methods are required that explicitly model aliasing. Two such methods will be presented in Section 3.

The second part of the super-resolution methods is the reconstruction, in which a high resolution image is reconstructed from the registered low resolution images. As described above, a frequency domain algorithm was presented by Tsai and Huang.¹ Kim et al.²¹ extended this algorithm to include blur and noise models. More recently several other, mostly spatial domain techniques have been developed. A good overview is given by Park et al.²² First of all, there are a set of nonuniform interpolation methods^{23,24} that reconstruct a bandlimited signal from an irregular set of samples. Another category of reconstruction algorithms is based on the POCS method.²⁵ These are iterative algorithms that project the reconstruction successively onto different convex sets representing prior information about the high resolution image. After a number of iterations, the reconstructed image converges to the image at the intersection of the different sets. Similarly, the iterative backprojection algorithms^{11,12,26} use a model of the imaging process to estimate the low resolution images that would be created when capturing the current reconstruction estimate. The high resolution estimate is then corrected with a term relative to the difference between these estimated low resolution images and the real ones. Finally, a maximum a posteriori (MAP) approach can also be used for reconstruction.^{9,27,28} It allows to easily add different priors on the reconstructed image into the model. Farsiu et al.²⁹ increased the robustness of the algorithm by Elad and Hel-Or²⁸ by replacing the L_2 minimization by an L_1 minimization. They also added a regularization term to stabilize the reconstruction.

This gives only a very concise overview of the wide variety of existing super-resolution reconstruction algorithms. For a more detailed overview, the reader is referred to the review articles mentioned above. We will concentrate here on the first part of most super-resolution algorithms, the image registration. More specifically, we will present some new registration algorithms for undersampled images.

3. METHODS

In this section, we will present three registration methods for aliased images. To simplify notations, the methods will be described for one-dimensional signals. They can straightforwardly be extended to two dimensions. Only planar shifts will be considered, because they are easy to describe in Fourier domain. First, we will give a brief analysis of aliasing (Section 3.1), and set up the mathematical framework. In Section 3.2, we present a method for registration of images that are partially aliased. And finally, in Section 3.3 and 3.4, we present two methods to register highly undersampled images.

3.1. Aliasing

Assume $f(t)$ is a bandlimited, periodic signal with maximum frequency L and period $T = 1$. Its Fourier expansion coefficients are denoted by α_l , with $-L \leq l \leq L$. We sample $f(t)$ at a frequency N , resulting in the samples

$$y_n = f(n/N), \quad \text{with } 0 \leq n < N. \quad (1)$$

If the sampling frequency satisfies the Shannon sampling theorem, or $N > 2L$, the samples y_n and their discrete Fourier transform coefficients Y_p can be written as

$$y_n = \sum_{l=-L}^L e^{j2\pi ln/N} \alpha_l, \quad \text{with } 0 \leq n < N$$

$$Y_p = \frac{1}{N} \sum_{n=0}^{N-1} e^{-j2\pi pn/N} y_n = \frac{1}{N} \sum_{l=-L}^L \alpha_l \sum_{n=0}^{N-1} e^{j2\pi(l-p)n/N} = \alpha_p, \quad \text{with } 0 \leq p < N. \quad (2)$$

However, if $N \leq 2L$, the signal is undersampled, and instead of (2), we obtain

$$y_n = \sum_{l=-L}^L e^{j2\pi ln/N} \alpha_l, \quad \text{with } 0 \leq n < N$$

$$Y_p = \frac{1}{N} \sum_{n=0}^{N-1} e^{-j2\pi pn/N} y_n = \frac{1}{N} \sum_{l=-L}^L \alpha_l \sum_{n=0}^{N-1} e^{j2\pi(l-p)n/N} = \sum_i \alpha_{p+iN}, \quad \text{with } 0 \leq p < N, \quad (3)$$

where the index i goes over all the nonzero overlapping spectrum copies. The signal coefficients α_l cannot be uniquely determined from a single set of samples y anymore. Other, lower frequency signals also fit through the same set of sampling locations. In images, this is visible as low frequency patterns that appear with different directions from the original pattern, or at edges that get a 'staircase' effect. Illustrations of these artifacts are given in Figure 1.



Figure 1. Illustrations of aliased images. Note the ringing artifacts on the walls in (a) and the guitar strings in (b).

Similarly, if we have two sets of samples y^0 and y^1 , taken with relative offset t_1 , we can write the samples y_n^1 and the Fourier transform coefficients Y_p^1 from this second set as

$$\begin{aligned}
 y_n^1 &= f\left(\frac{n}{N} + t_1\right) = \sum_{l=-L}^L e^{j2\pi l(n/N+t_1)} \alpha_l \\
 Y_p^1 &= \frac{1}{N} \sum_{l=-L}^L e^{j2\pi l t_1} \alpha_l \sum_{n=0}^{N-1} e^{j2\pi(l-p)n/N}
 \end{aligned} \tag{4}$$

If $N > 2L$, the offset t_1 can be easily determined using

$$\frac{Y_p^1}{Y_p^0} = \frac{\alpha_p e^{j2\pi p t_1}}{\alpha_p} = e^{j2\pi p t_1}, \tag{5}$$

with Y_p^0 and Y_p^1 the Fourier transforms of y^0 and y^1 , respectively. However, if $N \leq 2L$, we have

$$\frac{Y_p^1}{Y_p^0} = \frac{\sum_i \alpha_{p+iN} e^{j2\pi(p+iN)t_1}}{\sum_i \alpha_{p+iN}}, \tag{6}$$

and the offset cannot be directly retrieved anymore.

3.2. Registration of partially aliased images

If only part of the frequency spectrum is aliased, or $L < N \leq 2L$, the registration can still be performed using the aliasing-free part of the spectrum.²⁰ The frequencies l with $L - N < l < N - L$ are not aliased, and (5) still holds. The shift parameters can therefore be estimated as the linear phase difference between the Fourier transforms Y^0 and Y^1 of the sampled signals for the frequencies $L - N < l < N - L$.

3.3. SVD-based registration of completely aliased images

The approach described above is not valid anymore when $N < L$. In that case, the whole frequency spectrum is aliased, and (6) has to be used at any frequency. The aliasing has to be added to the signal model in order to be able to register the different signals.³⁰

Let us call the undersampling factor S ($S = \lceil 2L/N \rceil$). We will take S to be odd to simplify notations. Assume there are M sets of samples available: y^0, y^1, \dots, y^{M-1} . For accurate registration and reconstruction, we need at least $M > S$ sets of samples*.

The Fourier transform coefficients of each of the sampled sets can be written as

$$Y_p^m = \sum_{i=-(S-1)/2}^{(S-1)/2} e^{j2\pi t_m(p+iN)} \alpha_{p+iN} = e^{j2\pi t_m p} \sum_{i=-(S-1)/2}^{(S-1)/2} e^{j2\pi t_m iN} \alpha_{p+iN}, \quad (7)$$

with Y_p^m the p -th Fourier coefficient of the m -th sampling set y^m and t_m the offset of y^m with respect to y^0 . After putting the coefficients Y_p^m (with $0 \leq p < N$) into a vector, we obtain

$$\begin{pmatrix} Y_0^m \\ Y_1^m \\ \vdots \\ Y_{N-1}^m \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{j2\pi t_m} & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & e^{j2\pi(N-1)t_m} \end{pmatrix} \sum_{i=-(S-1)/2}^{(S-1)/2} e^{j2\pi t_m iN} \begin{pmatrix} \alpha_{iN} \\ \alpha_{iN+1} \\ \vdots \\ \alpha_{(i+1)N-1} \end{pmatrix}, \quad (8)$$

or

$$\mathbf{Y}^m = \mathbf{D}^{t_m} \sum_{i=-(S-1)/2}^{(S-1)/2} e^{j2\pi t_m iN} \boldsymbol{\alpha}^i, \quad (9)$$

where $\boldsymbol{\alpha}^i$ is the vector containing signal expansion coefficients α_{iN} to $\alpha_{(i+1)N-1}$. If we pre-multiply both sides of (9) by $(\mathbf{D}^{t_m})^{-1}$, we get

$$(\mathbf{D}^{t_m})^{-1} \mathbf{Y}^m = \sum_{i=-(S-1)/2}^{(S-1)/2} e^{j2\pi t_m iN} \boldsymbol{\alpha}^i. \quad (10)$$

From (10), we can see that the modified Fourier transforms of the sample vectors $(\mathbf{D}^{t_m})^{-1} \mathbf{Y}^m$ ($0 \leq m < M$) all belong to the same S -dimensional space $\text{span}(\boldsymbol{\alpha}^{-(S-1)/2}, \dots, \boldsymbol{\alpha}^{(S-1)/2})$. However, (10) is only valid for the correct offset values $\mathbf{t} = (t_0 \ t_1 \ \dots \ t_{M-1})$. For any other set of offset values $\hat{\mathbf{t}} = (\hat{t}_0 \ \hat{t}_1 \ \dots \ \hat{t}_{M-1})$, the space spanned by the M vectors $(\mathbf{D}^{\hat{t}_m})^{-1} \mathbf{Y}^m$ is M -dimensional.

The offsets \mathbf{t} can therefore be searched as the values for which the rank of the matrix

$$\begin{pmatrix} \mathbf{Y}^0 & (\mathbf{D}^{\hat{t}_1})^{-1} \mathbf{Y}^1 & \dots & (\mathbf{D}^{\hat{t}_{M-1}})^{-1} \mathbf{Y}^{M-1} \end{pmatrix} = \begin{pmatrix} Y_0^0 & Y_0^1 & \dots & Y_0^{M-1} \\ Y_1^0 & e^{j2\pi \hat{t}_1} Y_1^1 & \dots & e^{j2\pi \hat{t}_{M-1}} Y_1^{M-1} \\ Y_2^0 & e^{j4\pi \hat{t}_1} Y_2^1 & \dots & e^{j4\pi \hat{t}_{M-1}} Y_2^{M-1} \\ \vdots & \vdots & \dots & \vdots \\ Y_{N-1}^0 & e^{j2(N-1)\pi \hat{t}_1} Y_{N-1}^1 & \dots & e^{j2(N-1)\pi \hat{t}_{M-1}} Y_{N-1}^{M-1} \end{pmatrix} \quad (11)$$

*For images, this translates into $M > S^2$ input images, as we want to increase the resolution in both dimensions.

is S . Equivalently, we can search for the offset values $\hat{\mathbf{t}}$ such that the $S + 1$ -th singular value of (11) is minimized:

$$\min_{\hat{\mathbf{t}}} \sigma_{S+1}(\mathbf{Y}^{\hat{\mathbf{t}}}), \quad (12)$$

where $\mathbf{Y}^{\hat{\mathbf{t}}}$ is the matrix from (11), and $\sigma_{S+1}(\mathbf{Y}^{\hat{\mathbf{t}}})$ represents the $S + 1$ -th singular value of the matrix $\mathbf{Y}^{\hat{\mathbf{t}}}$.

Note that, as opposed to the method from Section 3.2 and most other image registration methods, all the M sets of samples have to be registered jointly (or at least $S + 1$ of them). Due to the undersampling, pairwise alignment of the different sets is not possible anymore. Tsai and Huang¹ also presented such a joint approach to registration.

3.4. Projection-based registration of completely aliased images

Another registration method for aliased signals (with $N < L$) is based on projections onto subspaces.³¹ Using (3), we can write

$$\begin{pmatrix} y_0^m \\ y_1^m \\ y_2^m \\ \vdots \\ y_{N-1}^m \end{pmatrix} = \begin{pmatrix} 1 & \dots & 1 & 1 & 1 & \dots & 1 \\ e^{-j2\pi Lt_m} e^{-j2\pi L/N} & \dots & e^{-j2\pi t_m} e^{-j2\pi/N} & 1 & e^{j2\pi t_m} e^{j2\pi/N} & \dots & e^{-j2\pi Lt_m} e^{j2\pi L/N} \\ e^{-j2\pi Lt_m} e^{-j2\pi 2L/N} & \dots & e^{-j2\pi t_m} e^{-j2\pi 2/N} & 1 & e^{j2\pi t_m} e^{j2\pi 2/N} & \dots & e^{j2\pi Lt_m} e^{j2\pi 2L/N} \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ e^{-j2\pi Lt_m} e^{-j2\pi(N-1)L/N} & \dots & e^{-j2\pi t_m} e^{-j2\pi(N-1)/N} & 1 & e^{j2\pi t_m} e^{j2\pi(N-1)/N} & \dots & e^{j2\pi Lt_m} e^{j2\pi(N-1)L/N} \end{pmatrix} \begin{pmatrix} \alpha_{-L} \\ \vdots \\ \alpha_{-1} \\ \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_L \end{pmatrix}, \quad (13)$$

or

$$\mathbf{y}^m = \Phi^{t_m} \boldsymbol{\alpha}, \quad (14)$$

where \mathbf{y}^m is the m -th sample vector, Φ^{t_m} is the shifted Fourier basis matrix, and $\boldsymbol{\alpha}$ is the vector containing the signal expansion coefficients α_l . Combining the different sample vectors into a large vector \mathbf{y} , we obtain

$$\mathbf{y} = \begin{pmatrix} \mathbf{y}^0 \\ \mathbf{y}^1 \\ \vdots \\ \mathbf{y}^{M-1} \end{pmatrix} = \begin{pmatrix} \Phi^{t_0} \\ \Phi^{t_1} \\ \vdots \\ \Phi^{t_{M-1}} \end{pmatrix} \boldsymbol{\alpha} = \Phi^{\mathbf{t}} \boldsymbol{\alpha}. \quad (15)$$

This shows that for the correct offsets \mathbf{t} , the sample vector \mathbf{y} is part of the space spanned by the columns of $\Phi^{\mathbf{t}}$, which is not the case for arbitrary offsets $\hat{\mathbf{t}}$.

It is therefore only for the correct values of \mathbf{t} that the projection of \mathbf{y} onto the space spanned by the columns of $\Phi^{\mathbf{t}}$ satisfies

$$\hat{\mathbf{y}} = \mathbf{P}_{\Phi^{\mathbf{t}}} \mathbf{y} = \Phi^{\mathbf{t}} (\Phi^{\mathbf{t}T} \Phi^{\mathbf{t}})^{-1} \Phi^{\mathbf{t}T} \mathbf{y} = \mathbf{y}. \quad (16)$$

The offsets can then be computed as the values of \mathbf{t} that minimize

$$\min_{\hat{\mathbf{t}}} \|\hat{\mathbf{y}} - \mathbf{y}\|^2 = \min_{\hat{\mathbf{t}}} \|\Phi^{\hat{\mathbf{t}}} (\Phi^{\hat{\mathbf{t}T} \Phi^{\hat{\mathbf{t}}})^{-1} \Phi^{\hat{\mathbf{t}T} \mathbf{y}} - \mathbf{y}\|^2. \quad (17)$$

Except for the columns that correspond to overlapping Fourier coefficients, all the columns of $\Phi^{\hat{\mathbf{t}}}$ are orthogonal to each other, for any arbitrary value of $\hat{\mathbf{t}}$. The minimization function from (17) can therefore be split according to the different orthogonal subspaces corresponding to sets of overlapping Fourier coefficients:

$$\min_{\hat{\mathbf{t}}} \|\hat{\mathbf{y}} - \mathbf{y}\|^2 = \min_{\hat{\mathbf{t}}} \sum_{n=0}^{N-1} \|\hat{\mathbf{y}}^{(n)} - \mathbf{y}^{(n)}\|^2, \quad (18)$$

with $\hat{\mathbf{y}}^{(n)}$ and $\mathbf{y}^{(n)}$ the projections onto the n -th subspace of $\hat{\mathbf{y}}$ and \mathbf{y} , respectively. In Figure 2, this decomposition is shown for two sets of samples from a one-dimensional signal. The minimization function is decomposed in four components, corresponding to the four orthogonal subspaces. The global minimum of (18) is located where the minima of the different components coincide. Such a decomposition reduces the computational complexity of the method significantly.

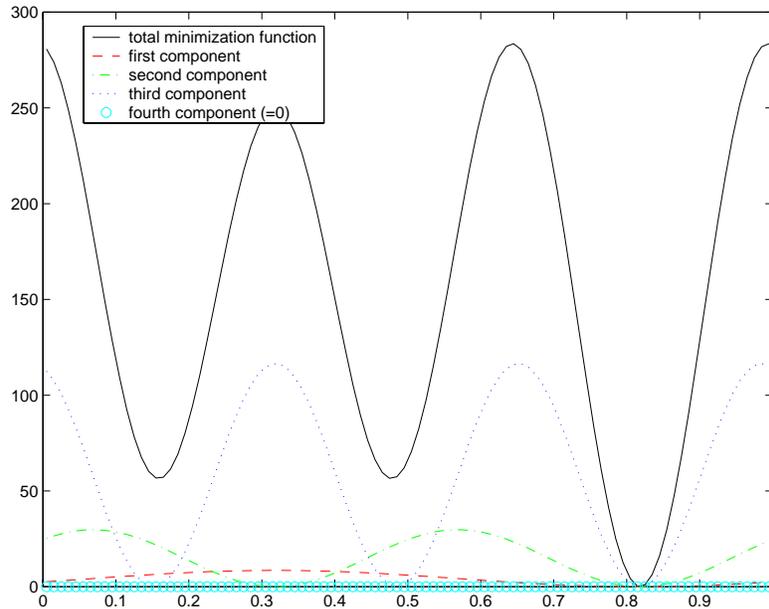


Figure 2. Example of the decomposition of the minimization function from (18) in its different components corresponding to orthogonal subspaces.

4. RESULTS

The registration methods described in Section 3 were tested in both simulations and practical experiments. First, the method based on an aliasing-free part of the spectrum (Section 3.2) is tested in a practical experiment. Four pictures of a real outdoor scene were taken using a Sigma SD10 digital camera. The images were captured while the camera was held manually in approximately the same position. This easily causes shifts of about ten pixels. The images clearly show aliasing artifacts (see Figure 3(a)). As small rotations are present in these images, the rotation estimation from Vandewalle et al.²⁰ was also applied to the aliasing-free frequencies of these images. For the reconstruction, we used a cubic interpolation method on the irregular set of samples. A high resolution image is reconstructed with twice the number of samples in both dimensions (Figure 3(b)). The reconstruction clearly shows more detail, and all aliasing has been accurately removed.

Next, the subspace methods from Section 3.3 and 3.4 were tested in simulations. A high resolution image was shifted using subpixel shifts and subsampled by slightly more than two, to generate five low resolution input images. The subsampling factor has to be slightly higher than the critical limit for reconstruction (which would be two in this case), to allow for the image registration. From these five images, the motion parameters were correctly estimated with both methods, and the original image is perfectly reconstructed (see Figure 4). Only the result using the SVD based registration method is shown, because the result with the projection method is the same up to computer working precision. Due to the computational complexity of these algorithms (see Section 5), the simulations are only performed for relatively small (128x128) images, and with a subsampling factor of two. For the image reconstruction in these simulations, the set of linear equations in (15) is solved using a least squares method.

The performance of the two subspace-based methods was also tested under noisy conditions. A one-dimensional signal is reconstructed from three aliased sets of samples for different amounts of white Gaussian noise. In this simulation, $2L + 1$ white Gaussian random variables were used as the Fourier coefficients of the (bandlimited) signal. The time domain signal was obtained using an inverse Fourier transform, and the different sets of samples are computed by sampling this signal with different offsets. Noise is then added to each of the sets of samples. The results are shown in Figure 5, and show the robustness of our methods to noise. Subpixel precision is obtained for SNRs above 20 dB.

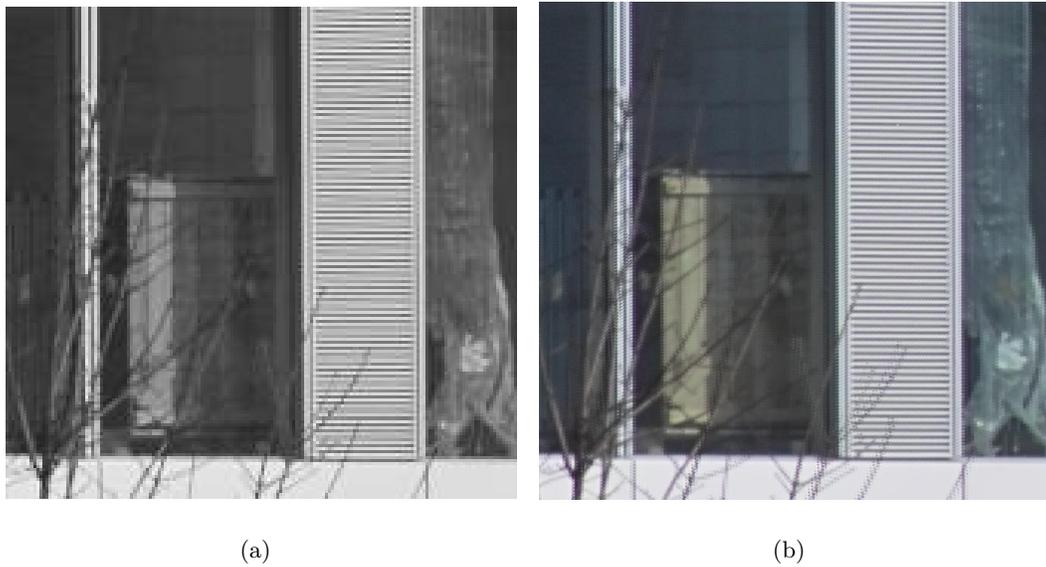


Figure 3. Results of the algorithm using an aliasing-free part of the spectrum in a practical experiment. (a) Detail from one of the four aliased input images to the super-resolution algorithm. (b) Detail of the resulting double resolution image.

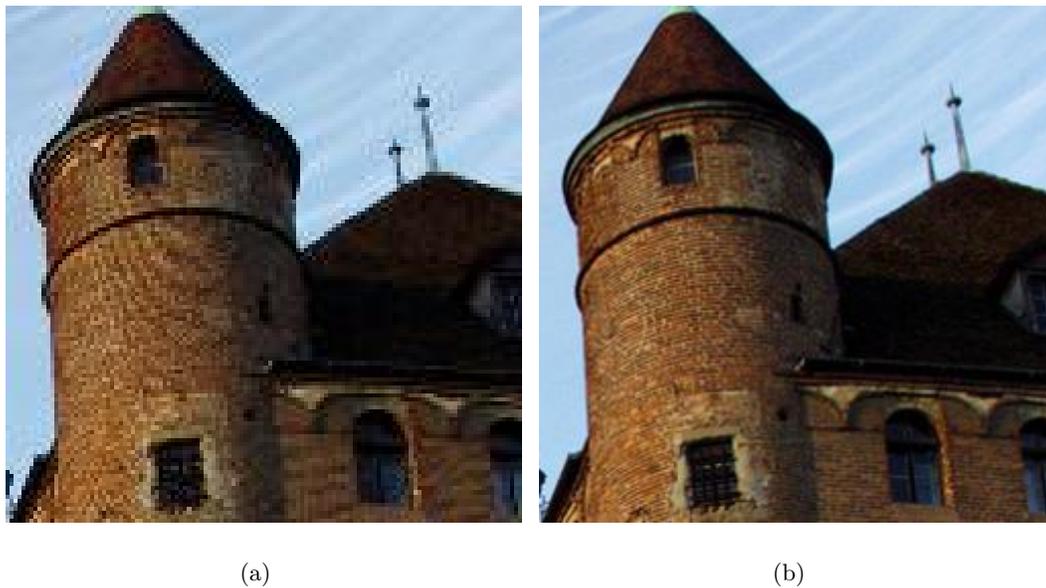


Figure 4. Results of the subspace-based algorithms. (a) One of the five input images to the algorithm (128×128). (b) The original high resolution (253×253) image is perfectly reconstructed from the five low resolution shifted images. Both the SVD-based algorithm from Section 3.3 and the projection-based algorithm from Section 3.4 give the same result.

5. DISCUSSION

The different methods described above all rely on the presence of aliasing for the reconstruction of a high resolution image. If the high frequency content of the images is filtered out before sampling, these methods do not give an increase in resolving power. This is the case for many consumer digital cameras. Many manufacturers consider the appearance of aliasing artifacts more disturbing than an overall blur of the image, and thus design

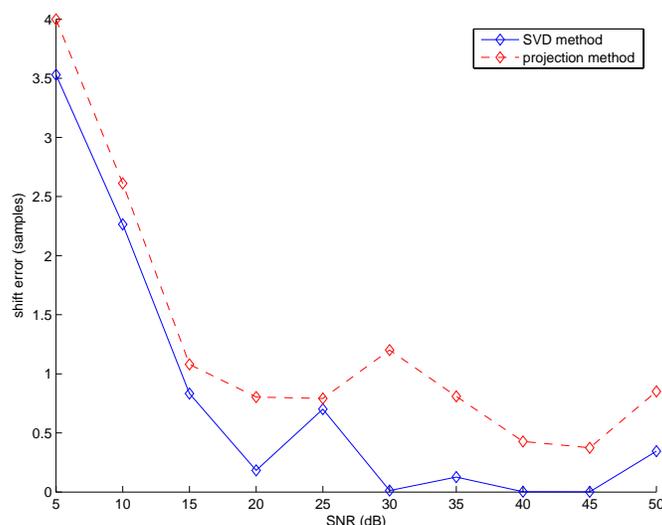


Figure 5. Results of the subspace-based algorithms on a 1D signal for different SNR values. 3 sets of 41 samples were used to reconstruct the 81 expansion coefficients. 50 simulations on random signals were performed.

the optical system accordingly.

No point spread function was considered in this article. In other words, the sampling is assumed to be Dirac sampling, and no blur is taken into account. However, as long as the point spread function is the same over the whole image, and if it does not remove certain frequencies entirely, this can easily be done after the reconstruction step. It is shown also by Elad and Hel-Or²⁸ that the optimality of the reconstruction is not modified by this separation between reconstruction and deblurring. The expansion coefficients α can be considered to be the product of the true expansion coefficients with the spatial frequency response of the camera. The coefficients can then be divided again by this spatial frequency response after reconstruction. Of course, if certain frequencies are highly attenuated, this could amplify the noise, thus degrading the reconstruction.

One of the main difficulties with the methods for completely aliased signals from sections 3.3 and 3.4 is their computational complexity. Both are M -dimensional minimization functions, with M the number of images. And due to the undersampling, this minimization has to be performed jointly in all dimensions. Moreover, as a function of the offsets $\hat{\mathbf{t}}$, these functions have many local minima, next to the global minimum (see Figure 6 for examples). If no prior estimate of the offsets is available, the function therefore has to be evaluated first on a regular M -dimensional grid of positions. Once a good approximation is obtained from this procedure, a standard minimization algorithm can be used to improve this initial estimate. However, in most cases we know a priori that the shift is in a certain range (for example thanks to a prior pixel-level registration, or by knowing that the shift is not larger than x pixels), and can thus avoid the first step.

As it describes the signal in function of an arbitrary basis matrix, the method from Section 3.4 is actually not limited to bandlimited images. It can be straightforwardly extended to other bases, like wavelets, splines, etc.

6. CONCLUSIONS

In this paper, we presented three frequency domain methods for the registration of undersampled images. First, a method was presented to align images that have a part of their spectrum free of aliasing. The performance of this method is shown in a practical experiment. Next, two methods are described to align images for which the entire frequency spectrum is aliased. They model also the aliased part of the spectrum, and use a subspace method to find the offset parameters. The main issue with these two algorithms is their high computational complexity. They are tested in noiseless and noisy conditions in simulations and showed good results.

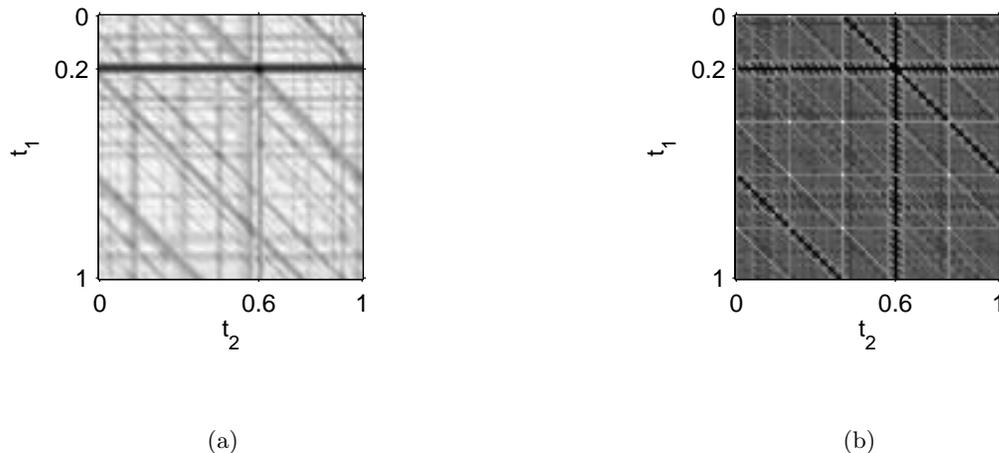


Figure 6. Local minima in the minimization function from (12) (left) and (17) (right). Higher values are represented by brighter grayscale values. For both images, three sets of 1D functions were used, with correct offset values $t_1 = 0.2$ and $t_2 = 0.6$.

ACKNOWLEDGMENTS

We would like to thank Foveon, Inc. for providing a Sigma SD10 digital camera and the image in Figure 1(b). The work presented in this paper was supported by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS, <http://www.mics.org>), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

REFERENCES

1. R. Y. Tsai and T. S. Huang, "Multiframe image restoration and registration," in *Advances in Computer Vision and Image Processing*, T. S. Huang, ed., 1, 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. "IEEE Signal Processing Magazine, special issue on super-resolution," May 2003.
4. "EURASIP journal on applied signal processing, special issue on super-resolution," 2005.
5. S. Farsiu, D. Robinson, and P. Milanfar, "MDSP resolution enhancement software.", [Online]. Available: <http://www.soe.ucsc.edu/milanfar/SR-Software.htm>, 2004.
6. P. Vandewalle, P. Zbinden, S. Süsstrunk, and M. Vetterli, "Super-resolution software.", [Online]. Available: <http://lcavwww.epfl.ch/software/superresolution>, 2005.
7. L. G. Brown, "A survey of image registration techniques," *ACM Comput. Surv.* **24**(4), pp. 325–376, 1992.
8. B. Zitová and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing* **21**(11), pp. 977–1000, 2003.
9. D. Capel and A. Zisserman, "Computer vision applied to super resolution," *IEEE Signal Processing Magazine* **20**, pp. 75–86, May 2003.
10. M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM* **24**, pp. 381–395, June 1981.
11. M. Irani and S. Peleg, "Improving resolution by image registration," *CVGIP: Graphical Models and Image Processing* **53**, pp. 231–239, May 1991.
12. 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.

13. B. Galvin, B. McCane, K. Novins, D. Mason, and S. Mills, "Recovering motion fields: An evaluation of eight optical flow algorithms," in *Proceedings of the British Machine Vision Conference*, pp. 454–460, September 1998.
14. H. Foroosh, J. B. Zerubia, and M. Berthod, "Extension of Phase Correlation to Subpixel Registration," *IEEE Transactions on Image Processing* **11**(3), pp. 188–200, 2002.
15. B. S. Reddy and B. N. Chatterji, "An fft-based technique for translation, rotation and scale-invariant image registration," *IEEE Transactions on Image Processing* **5**, pp. 1266–1271, August 1996.
16. 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.
17. L. Lucchese and G. M. Cortelazzo, "A noise-robust frequency domain technique for estimating planar roto-translations," *IEEE Transactions on Signal Processing* **48**, pp. 1769–1786, June 2000.
18. 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.
19. 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.
20. P. Vandewalle, S. Süssstrunk, and M. Vetterli, "A frequency domain approach to registration of aliased images with application to super-resolution," accepted to *EURASIP Journal on Applied Signal Processing, Special Issue on Super-Resolution Imaging*, 2005.
21. S. P. Kim, N. K. Bose, and H. M. Valenzuela, "Recursive reconstruction of high resolution image from noisy undersampled multiframes," *IEEE Transactions on Acoustics, Speech, and Signal Processing* **38**, pp. 1013–1027, June 1990.
22. 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.
23. A. Papoulis, "Generalized sampling expansion," *IEEE Transactions on Circuits and Systems* **24**, pp. 652–654, November 1977.
24. T. Strohmer, "Computationally attractive reconstruction of bandlimited images from irregular samples," *IEEE Transactions on Image Processing* **6**, pp. 540–548, April 1997.
25. A. J. Patti, M. I. Sezan, and A. M. Tekalp, "Superresolution video reconstruction with arbitrary sampling lattices and nonzero aperture time," *IEEE Transactions on Image Processing* **6**, pp. 1064–1076, August 1997.
26. B. C. Tom and A. K. Katsaggelos, "Resolution enhancement of monochrome and color video using motion compensation," *IEEE Transactions on Image Processing* **10**, pp. 278–287, February 2001.
27. R. R. Schultz, L. Meng, and R. L. Stevenson, "Subpixel motion estimation for super-resolution image sequence enhancement," *Journal of Visual Communication and Image Representation* **9**, pp. 38–50, March 1998.
28. M. Elad and Y. Hel-Or, "A Fast Super-Resolution Reconstruction Algorithm for Pure Translational Motion and Common Space-Invariant Blur," *IEEE Transactions on Image Processing* **10**, pp. 1187–1193, Aug. 2001.
29. S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, "Fast and robust multiframe super resolution," *IEEE Transactions on Image Processing* **13**, pp. 1327–1344, October 2004.
30. P. Vandewalle, L. Sbaiz, J. Vandewalle, and M. Vetterli, "How to take advantage of aliasing in bandlimited signals," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, **3**, pp. 948–951, May 2004.
31. P. Vandewalle, L. Sbaiz, M. Vetterli, and S. Süssstrunk, "Super-resolution from highly undersampled images," in *IEEE International Conference on Image Processing*, September 2005.