

A Frequency Domain Approach to Registration of Aliased Images with Application to Super-resolution

Patrick Vandewalle,¹ Sabine Süsstrunk,¹ and Martin Vetterli^{1,2}

¹*Ecole Polytechnique Fédérale de Lausanne, School of Computer and Communication Sciences, 1015 Lausanne, Switzerland*

²*Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720-1770, USA*

Received 27 November 2004; Revised 4 May 2005; Accepted 18 May 2005

Super-resolution algorithms reconstruct a high-resolution image from a set of low-resolution images of a scene. Precise alignment of the input images is an essential part of such algorithms. If the low-resolution images are undersampled and have aliasing artifacts, the performance of standard registration algorithms decreases. We propose a frequency domain technique to precisely register a set of aliased images, based on their low-frequency, aliasing-free part. A high-resolution image is then reconstructed using cubic interpolation. Our algorithm is compared to other algorithms in simulations and practical experiments using real aliased images. Both show very good visual results and prove the attractiveness of our approach in the case of aliased input images. A possible application is to digital cameras where a set of rapidly acquired images can be used to recover a higher-resolution final image.

Copyright © 2006 Patrick Vandewalle et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Image resolution is one of the limiting parameters in digital camera design. With most digital cameras, however, it is possible to take bursts of multiple pictures in a very short period of time. Thus, high-resolution images can be reconstructed from a series of low-resolution images using super-resolution algorithms.

The idea behind super-resolution imaging is to combine the information from a set of slightly different low-resolution images of the same scene and use it to construct a higher-resolution image. Throughout this paper, a higher-resolution image is defined as an image with more resolving power. This means that an image that is obtained by merely upsampling and interpolating a low-resolution image does not have a higher resolution than its original. It has a larger number of pixels, but the resolving power remains the same; that is, the interpolated image does not contain more details than its original. The resolving power of an image can be increased by adding high-frequency information typically based on knowledge about the specific image model. A higher resolving power is also obtained when the aliasing ambiguity in an image is removed. We will take this second approach to construct high-resolution images. The aliasing ambiguity in an image is removed by incorporating the additional information obtained from other images of the same scene.

There are two major, and to some extent, independent challenges in super-resolution imaging. First, the difference between the low-resolution input images needs to be known precisely. This difference can have many origins: camera motion [1–7], change of focus [8, 9], a combination of these two [10–13], and so forth. We will consider images that differ by a planar motion. Therefore, the first challenge corresponds to having a precise knowledge of the motion parameters. This is a challenge because we use images containing possibly large amounts of aliasing. An error in the motion estimation translates almost directly into a degradation of the resulting high-resolution image. It is generally better to interpolate one of the low-resolution images than to create a high-resolution image from the set of images using incorrect motion parameters. The artifacts caused by an incorrectly aligned image are visually much more disturbing than the blurring effect from interpolating only one image. The second challenge is to apply the information obtained from the different registered images to the reconstruction of a sharp high-resolution image. A nontrivial deconvolution operation is required to undo the blurring operation applied by the camera point spread function. This paper mainly focuses on the first problem, thus, no point spread function is taken into account in the reconstruction. The sampling operation is assumed to be ideal Dirac sampling and we do not consider the deconvolution problem.

We describe an image registration algorithm using a new frequency domain method that outperforms the state of the art in frequency domain registration methods. It also performs better than the spatial domain method by Keren et al. [5] if the images have some directionality. Unlike other motion estimation algorithms whose performance is often very low for noisy or highly aliased images (see Section 2), our algorithm only uses low-frequency information. This is the part of the signals with the highest signal-to-noise ratio (SNR), and in our setup, the aliasing-free part of the images. We developed a new, computationally efficient method to estimate planar rotations. To reconstruct the high-resolution image, we apply bicubic interpolation on a high-resolution grid. The super-resolution algorithm we propose reconstructs an image with almost double resolution in both dimensions from four aliased images. The four low-resolution images are necessarily undersampled. Otherwise, our algorithm is not able to reconstruct a better image as it uses exactly this undersampled information. We compare our approach in a simulation to other spatial domain and frequency domain registration algorithms. We find that our algorithm can better estimate shift and rotation parameters than the other methods, in particular, when some strong directionality is present in the image.

A possible application of the proposed image registration algorithm is that of a user holding his digital camera in his hands while manually or automatically taking a series of four shots of a scene within a short period of time. The small vibrations of the user's hands during image capture are sufficient to reconstruct a high-resolution image. The scene needs to be flat or at a large distance, such that no parallax effects take place. We tested such a setup using real digital cameras. We verified that aliasing occurs with these cameras by measuring their spatial frequency response. In the experiments, we found that our algorithm results in better visual quality than the other methods, which typically failed to adequately register all four images. Other applications of super-resolution algorithms can be found in forensic imaging, satellite imaging, microscopy, medical imaging, constructing still images from video sequences, and so forth.

The article is organized as follows. Section 2 discusses the state of the art in image registration and super-resolution imaging. The planar motion estimation algorithm is described in Section 3 and the reconstruction in Section 4. Section 5 shows the results on simulated and real images and the comparison to other algorithms. The results are discussed in Sections 6 and 7 concluding the article.

2. STATE OF THE ART

The idea of super-resolution was first introduced in 1984 by Tsai and Huang [1] for multiframe image restoration of bandlimited signals. A good overview of existing algorithms is given by Borman and Stevenson [14] and Park et al. [15]. Most super-resolution methods are composed of two main steps: first all the images are aligned in the same coordinate system in the registration step, and then a high-resolution image is reconstructed from the irregular set of samples. In

this second step, the camera point spread function is often taken into account.

Precise subpixel image registration is a basic requirement for a good reconstruction. If the images are inaccurately registered, the high-resolution image is reconstructed from incorrect data and is not a good representation of the original signal. Zitová and Flusser [16] present an overview of image registration methods. Registration can be done either in spatial or in frequency domain. By the nature of the Fourier transform, frequency domain methods are limited to global motion models. In general, they also consider only planar shifts and possibly planar rotation and scale, which can be easily expressed in Fourier domain. However, aliasing is much easier to describe and to handle in frequency domain than in spatial domain.

Tsai and Huang [1] describe an algorithm to register multiple frames simultaneously using nonlinear minimization in frequency domain. Their method for registering multiple aliased images is based on the fact that the original, high-resolution signal is bandlimited. It is not clear, however, if such a solution is unique and if such an algorithm will not converge to a local minimum. Most of the frequency domain registration methods are based on the fact that two shifted images differ in frequency domain by a phase shift only, which can be found from their correlation. Using a log-polar transform of the magnitude of the frequency spectra, image rotation and scale can be converted into horizontal and vertical shifts. These can therefore also be estimated using a phase correlation method. Reddy and Chatterji [17] and Marcel et al. [18] describe such planar motion estimation algorithms. Reddy and Chatterji apply a high-pass emphasis filter to strengthen high frequencies in the estimation. Kim and Su [19], Stone et al. [20], and Vandewalle et al. [2] also apply a phase correlation technique to estimate planar shifts. To minimize errors due to aliasing, their methods rely on a part of the frequency spectrum that is almost free of aliasing. Typically this is the low-frequency part of the images. Foroosh et al. [21] showed that the signal power in the phase correlation corresponds to a polyphase transform of a filtered unit impulse. Lucchese and Cortelazzo [22] developed a rotation estimation algorithm based on the property that the magnitude of the Fourier transform of an image and the mirrored version of the magnitude of the Fourier transform of a rotated image have a pair of orthogonal zero-crossing lines. The angle that these lines make with the axes is equal to half the rotation angle between the two images. The horizontal and vertical shifts are estimated afterwards using a standard phase correlation method.

Spatial domain methods generally allow for more general motion models, such as homographies. They can be based on the whole image or on a set of selected corresponding feature vectors, as discussed by Capel and Zisserman [4] and by Fischler and Bolles in their RANSAC algorithm [23]. Keren et al. [5] developed an iterative planar motion estimation algorithm based on Taylor expansions. A pyramidal scheme is used to increase the precision for large motion parameters. Bergen et al. developed a hierarchical framework to estimate motion in a multiresolution data structure [24]. Different

motion models, such as affine flow or rigid body motion, can be used in combination with this approach. Irani et al. [25] present a method to compute multiple, possibly transparent or occluding motions in an image sequence. Motion is estimated using an iterative multiresolution approach based on planar motion. Different objects are tracked using segmentation and temporal integration. Gluckman [26] describes a method that first computes planar rotation from the gradient field distribution of the images to be registered. Planar shifts are then estimated after cancellation of the rotation using a phase correlation method.

In the subsequent image reconstruction phase, a high-resolution image is reconstructed from the irregular set of samples that is obtained from the different low-resolution images. This can be achieved using an interpolation-based method as the one used by Keren et al. [5]. Tsai and Huang [1] describe a frequency domain method, writing the Fourier coefficients of the high-resolution image as a function of the Fourier coefficients of the registered low-resolution images. The solution is then computed from a set of linear equations. This algorithm uses the same principle as the formulation in time domain given by Papoulis [27]. A high-resolution image can also be reconstructed using a POCS algorithm (Patti et al. [10]), where the estimated reconstruction is successively projected on different convex sets. Each set represents constraints to the reconstructed image that are based on the given measurements and assumptions about the signal. Capel and Zisserman [4] and Schultz et al. [6] use a maximum a posteriori (MAP) statistical method to build the high-resolution image.

Other methods iteratively create a set of low-resolution images from the estimated image using the imaging model. The estimate is then updated according to the difference between the real and the simulated low-resolution images (Keren et al. [5], Irani and Peleg [7]). This method is known as iterative backprojection. Zomet et al. [11] improved the results obtained with typical iterative backprojection algorithms by taking the median of the errors in the different backprojected images. This proved to be more robust in the presence of outliers. Farsiu et al. [12] proposed a new and robust super-resolution algorithm. Instead of the more common L_2 minimization, they use the L_1 norm, which produces sharper high-resolution images. They also showed that this approach performs very well in combination with the algorithm by Zomet et al. [11]. Elad and Feuer [13] present a super-resolution framework that combines a maximum-likelihood/MAP approach with a POCS approach to define a new convex optimization problem. Next, they show the connections between their method and different classes of other existing methods.

Our main contribution in this paper consists of a new frequency domain algorithm to register not just low resolution, but also aliased images. We use a planar motion model. When a series of images is taken in a short amount of time with only small camera motion between the images, we assume that the motion can be described with such a model. In general, a planar model is simpler and has less parameters making it often more robust in the presence of noise. We

also extend the planar shift motion model from [2, 19, 20] to include planar rotations, because they are often part of the camera motion. Even a small rotation has a large influence on final registration. Our rotation estimation algorithm is computationally efficient and adapted to work with aliased images. We test our algorithm not only in simulations, but also on real sequences of aliased images. The results from these tests validate the assumptions made about the motion. They show, both visually and in SNR, that our algorithm outperforms other frequency domain registration methods as well as a spatial domain method if directionality is present in the images.

3. PLANAR MOTION ESTIMATION

We use a frequency domain algorithm to estimate the motion parameters between the reference image and each of the other images. Only planar motion parallel to the image plane is allowed. 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 .

A frequency domain approach allows us to estimate the horizontal and vertical shift and the (planar) rotation separately. Assume we have a reference signal $f_1(\mathbf{x})$ and its shifted and rotated version $f_2(\mathbf{x})$:

$$f_2(\mathbf{x}) = f_1(\mathbf{R}(\mathbf{x} + \Delta\mathbf{x})),$$

$$\text{with } \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}, \quad \Delta\mathbf{x} = \begin{bmatrix} \Delta x_1 \\ \Delta x_2 \end{bmatrix}, \quad \mathbf{R} = \begin{bmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{bmatrix}. \quad (1)$$

This can be expressed in 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 $F_2(\mathbf{u})$ the Fourier transform of $f_2(\mathbf{x})$ and the coordinate transformation $\mathbf{x}' = \mathbf{x} + \Delta\mathbf{x}$. After another transformation $\mathbf{x}'' = \mathbf{R}\mathbf{x}'$, the relation between the amplitudes of the Fourier transforms can be computed as

$$\begin{aligned} |F_2(\mathbf{u})| &= \left| e^{j2\pi\mathbf{u}^T \Delta\mathbf{x}} \iint_{\mathbf{x}'} f_1(\mathbf{R}\mathbf{x}') e^{-j2\pi\mathbf{u}^T \mathbf{x}'} d\mathbf{x}' \right| \\ &= \left| \iint_{\mathbf{x}'} f_1(\mathbf{R}\mathbf{x}') e^{-j2\pi\mathbf{u}^T \mathbf{x}'} d\mathbf{x}' \right| \\ &= \left| \iint_{\mathbf{x}''} f_1(\mathbf{x}'') e^{-j2\pi\mathbf{u}^T (\mathbf{R}^T \mathbf{x}'')} d\mathbf{x}'' \right| \\ &= \left| \iint_{\mathbf{x}''} f_1(\mathbf{x}'') e^{-j2\pi(\mathbf{R}\mathbf{u})^T \mathbf{x}''} d\mathbf{x}'' \right| \\ &= |F_1(\mathbf{R}\mathbf{u})|, \end{aligned} \quad (3)$$




FIGURE 1: Rotation estimation. (a) Frequency values of the reference image for $0.1\rho < r < \rho$. (b) Frequency values of the rotated image ($\phi = 25$ degrees) for $0.1\rho < r < \rho$. (c) Average value as a function of the angle $h(\alpha)$ for both $|F_1(\mathbf{u})|$ and $|F_2(\mathbf{u})|$.

where $|F_2(\mathbf{u})|$ is a rotated version of $|F_1(\mathbf{u})|$ over the same angle ϕ as the spatial domain rotation (see Figures 1(a) and 1(b)). $|F_1(\mathbf{u})|$ and $|F_2(\mathbf{u})|$ do not depend on the shift values $\Delta\mathbf{x}$, because the spatial domain shifts only affect the phase values of the Fourier transforms. Therefore we can first estimate the rotation angle ϕ from the amplitudes of the Fourier transforms $|F_1(\mathbf{u})|$ and $|F_2(\mathbf{u})|$. After compensation for the rotation, the shift $\Delta\mathbf{x}$ can be computed from the phase difference between $F_1(\mathbf{u})$ and $F_2(\mathbf{u})$.

In Section 3.1, we give a precise rotation estimation algorithm. A subpixel shift estimation algorithm is described in Section 3.2, and an adaptation of this method to estimate motion accurately in aliased images is presented in Section 3.3.

3.1. Rotation estimation

The rotation angle between $|F_1(\mathbf{u})|$ and $|F_2(\mathbf{u})|$ can be computed as the angle θ for which the Fourier transform of the reference image $|F_1(\mathbf{u})|$ and the rotated Fourier transform of the image to be registered $|F_2(\mathbf{R}_\theta\mathbf{u})|$ have maximum correlation. This implies the computation of a rotation of $|F_2(\mathbf{u})|$ for every evaluation of the correlation, which is computationally heavy and thus practically difficult.

If $|F_1(\mathbf{u})|$ and $|F_2(\mathbf{u})|$ are transformed in polar coordinates, the rotation over the angle ϕ is reduced to a (circular) shift over ϕ . We can compute the Fourier transform of the spectra $|F_1(\mathbf{u})|$ and $|F_2(\mathbf{u})|$, and compute ϕ as the phase shift between the two (as it was also done by Marcel et al. [18] and Reddy and Chatterji [17]). This requires a transformation of the spectrum to polar coordinates. The data from the regular x_1, x_2 -grid need to be interpolated to obtain a regular r, θ -grid. Mainly for the low frequencies, which generally contain most of the energy, the interpolations are based on very few

function values and thus introduce large approximation errors. An implementation of this method is also computationally intensive.

Our approach is computationally much more efficient than the two methods described above. First of all, we compute the frequency content h as a function of the angle α by integrating over radial lines:

$$h(\alpha) = \int_{\alpha-\Delta\alpha/2}^{\alpha+\Delta\alpha/2} \int_0^\infty |F(r, \theta)| dr d\theta. \quad (4)$$

In practice, $|F(r, \theta)|$ is a discrete signal. Therefore, we compute the discrete function $h(\alpha)$ as the average of the values on the rectangular grid that have an angle $\alpha - \Delta\alpha/2 < \theta < \alpha + \Delta\alpha/2$. As we want to compute the rotation angle with a precision of 0.1 degrees, $h(\alpha)$ is computed every 0.1 degrees. To get a similar number of signal values $|F(r, \theta)|$ at every angle, the average is only evaluated on a circular disc of values for which $r < \rho$ (where ρ is the image radius or half the image size). Finally, as the values for low frequencies are very large compared to the other values and are very coarsely sampled as a function of the angle, we discard the values for which $r < \epsilon\rho$, with $\epsilon = 0.1$. Thus, $h(\alpha)$ is computed as the average of the frequency values on a discrete grid with $\alpha - \Delta\alpha/2 < \theta < \alpha + \Delta\alpha/2$ and $\epsilon\rho < r < \rho$.

This results in a function $h(\alpha)$ for both $|F_1(\mathbf{u})|$ and $|F_2(\mathbf{u})|$ (Figure 1(c)). The exact rotation angle can then be computed as the value for which their correlation reaches a maximum. Note that only a one-dimensional correlation has to be computed, as opposed to the two-dimensional correlation approaches presented in [17, 18].




FIGURE 2: In the presence of (partial) aliasing, the shift between two sampled signals cannot be found directly. However, after low-pass filtering, the shift can be easily determined. (a) Original continuous-time signal in time and frequency domain. (b) Sampled signal in time and frequency domain, with aliasing. (c) Low-pass filtered sampled signal in time and frequency domain.

3.2. Shift estimation

A shift of the image parallel to the image plane can be expressed in Fourier domain as a linear phase shift:

$$\begin{aligned} F_2(\mathbf{u}) &= \iint_{\mathbf{x}} f_2(\mathbf{x}) e^{-j2\pi\mathbf{u}^T \mathbf{x}} d\mathbf{x} = \iint_{\mathbf{x}} f_1(\mathbf{x} + \Delta\mathbf{x}) e^{-j2\pi\mathbf{u}^T \mathbf{x}} d\mathbf{x} \\ &= e^{j2\pi\mathbf{u}^T \Delta\mathbf{x}} \iint_{\mathbf{x}'} f_1(\mathbf{x}') e^{-j2\pi\mathbf{u}^T \mathbf{x}'} d\mathbf{x}' = e^{j2\pi\mathbf{u}^T \Delta\mathbf{x}} F_1(\mathbf{u}). \end{aligned} \quad (5)$$

It is well known that the shift parameters $\Delta\mathbf{x}$ can thus be computed as the slope of the phase difference $\angle(F_2(\mathbf{u})/F_1(\mathbf{u}))$. To make the solution less sensitive to noise, a plane is fitted through the phase differences using a least squares method.

3.3. Aliasing

If the low-resolution images are aliased, the methods described earlier do not result in precise registration anymore. This is due to the difference in frequency content of the low-resolution images caused by the aliasing. In this case, (2), (3), and (5) no longer hold. Instead of (5), a shift is now expressed as

$$F_2(\mathbf{u}) = \sum_{k=-K}^K e^{j2\pi(\mathbf{u}-k\mathbf{u}_s)^T \Delta\mathbf{x}} F_1(\mathbf{u} - k\mathbf{u}_s), \quad (6)$$

with \mathbf{u}_s the sampling frequency and $2K + 1$ overlapping spectrum copies at frequency \mathbf{u} . Aliasing terms disturb the linear phase relation between $F_1(\mathbf{u})$ and $F_2(\mathbf{u})$. However, in cases of limited aliasing, it is still possible to use the above methods, by considering only the frequencies that are free of aliasing or only marginally affected by aliasing. A similar idea was used for shift estimation methods by Kim and Su [19] and by Stone et al. [20].

Assume a one-dimensional, bandlimited signal $f(x)$ (with maximum frequency u_{\max} , Figure 2(a)), which is sampled at a frequency $u_{\max} < u_s < 2u_{\max}$. This does not satisfy the Nyquist criterion, and the sampled signal $f[k]$ will have aliasing artifacts (Figure 2(b)). $f(x)$ cannot be perfectly reconstructed from the samples $f[k]$. Consider two sampled signals, $f_1[k]$ and $f_2[k]$, sampled at $0, T, 2T, \dots, kT, \dots$ and $\Delta x, T + \Delta x, 2T + \Delta x, \dots, kT + \Delta x, \dots$, respectively (with $T = 1/u_s$ the sampling period). Due to the aliasing, their Fourier transforms differ by more than just a linear phase shift, and the shift estimation method described above does not work any more. However, the values at frequencies $-u_s + u_{\max} < u < u_s - u_{\max}$ are free of aliasing and thus the same for the two sampled signals $f_1[k]$ and $f_2[k]$ (up to a linear phase shift). So if a low-pass filter is applied to $f_1[k]$ and $f_2[k]$, the resulting signals $f_{1,\text{low}}[k]$ and $f_{2,\text{low}}[k]$ are exactly the same up to their shift Δx (Figure 2(c)). This shift can then be derived using a correlation operator in time domain or by estimating the linear phase difference in frequency domain.

An extension to two dimensions is straightforward. The two sampled signals $f_1[\mathbf{k}]$ and $f_2[\mathbf{k}]$ are first low-pass filtered (with cutoff frequency $\mathbf{u}_s - \mathbf{u}_{\max}$) in horizontal and vertical dimensions. The filtered images are identical up to their registration parameters and can be registered using the methods described in Sections 3.1 and 3.2. As both methods are applied in the Fourier domain, the filtering step can be avoided by applying the registration algorithms immediately to the low frequencies. The rotation estimation is then based on the frequencies for which $\epsilon\rho < r < \rho_{\max}$ (with $\rho_{\max} = \min((\mathbf{u}_s - \mathbf{u}_{\max})/\mathbf{u}_s)$), and the horizontal and vertical shifts are estimated from the phase differences for $-\mathbf{u}_s + \mathbf{u}_{\max} < \mathbf{u} < \mathbf{u}_s - \mathbf{u}_{\max}$.

Using this approach, high-frequency noise is removed together with the aliasing, which results in more accurate registration. A global overview of the registration algorithm is given in Algorithm 1.

- (1) Multiply the images $f_{LR,m}$ by a Tukey window to make them circularly symmetric. The windowed images are called $f_{LR,w,m}$.
- (2) Compute the Fourier transforms $F_{LR,w,m}$ of all low-resolution images.
- (3) *Rotation estimation*: the rotation angle between every image $f_{LR,w,m}$ ($m = 2, \dots, M$) and the reference image $f_{LR,w,1}$ is estimated.
 - (a) Compute the polar coordinates (r, θ) of the image samples.
 - (b) For every angle α , compute the average value $h_m(\alpha)$ of the Fourier coefficients for which $\alpha - 1 < \theta < \alpha + 1$ and $0.1\rho < r < \rho_{\max}$. The angles are expressed in degrees and $h_m(\alpha)$ is evaluated every 0.1 degrees. A typical value used for ρ_{\max} is 0.6.
 - (c) Find the maximum of the correlation between $h_1(\alpha)$ and $h_m(\alpha)$ between -30 and 30 degrees. This is the estimated rotation angle ϕ_m .
 - (d) Rotate image $f_{LR,w,m}$ by $-\phi_m$ to cancel the rotation.
- (4) *Shift estimation*: the horizontal and vertical shifts between every image $f_{LR,w,m}$ ($m = 2, \dots, M$) and the reference image $f_{LR,w,1}$ are estimated.
 - (a) Compute the phase difference between image m and the reference image as $\angle(F_{LR,w,m}/F_{LR,w,1})$.
 - (b) For all frequencies $-\mathbf{u}_s + \mathbf{u}_{\max} < \mathbf{u} < \mathbf{u}_s - \mathbf{u}_{\max}$ write the linear equation describing a plane through the computed phase difference with unknown slopes $\Delta\mathbf{x}$.
 - (c) Find the shift parameters $\Delta\mathbf{x}_m$ as the least squares solution of the equations.
- (5) *Image reconstruction*: a high-resolution image f_{HR} is reconstructed from the registered images $f_{LR,m}$ ($m = 1, \dots, M$).
 - (a) For every image $f_{LR,m}$, compute the coordinates of its pixels in the coordinate frame of $f_{LR,1}$ using the estimated registration parameters.
 - (b) From these known samples, interpolate the values on a regular high-resolution grid using for example cubic interpolation.

ALGORITHM 1: An overview of the complete super-resolution algorithm as it was described in Sections 3 and 4. A high-resolution image f_{HR} (with Fourier transform F_{HR}) is reconstructed from a set of M low-resolution images $f_{LR,m}$ ($m = 1, 2, \dots, M$) with Fourier transform $F_{LR,m}$.

4. RECONSTRUCTION

When the low-resolution images are accurately registered, the samples of the different images can be combined to reconstruct a high-resolution image. As discussed in Section 1, the sampling kernel is assumed to be a Dirac. In other words, no (generally low-pass filtering) point spread function was considered. For methods to deconvolve the image from a (known) point spread function, we refer to the reconstruction algorithms reviewed in Section 2.

In our reconstruction algorithm, the samples of the different low-resolution images are first expressed in the coordinate frame of the reference image. Then, based on these known samples, the image values are interpolated on a regular high-resolution grid. We chose bicubic interpolation because of its low computational complexity and good results.

What is the optimal number of images to use when reconstructing a high-resolution image? The exact answer to this question depends on many parameters, such as the registration accuracy, imaging model, total frequency content, and so forth. Intuitively, two effects need to be balanced. On one hand, the more images there are, the better the reconstruction should be. On the other hand, there is a limit to the improvements that can be obtained: even from a very large number of very low-resolution images of a scene, it will not be possible to reconstruct a sharp, high-resolution image. Blur, noise, and inaccuracies in the signal model limit the increase in resolving power that can be obtained. In our case, the motion estimation algorithm is limited to subsampling by a factor less than two in both dimensions (because our algorithm needs an aliasing-free part of the spectrum, see

also Section 3.3). Therefore, the resolution can only be really increased by (almost) a factor of four. Any supplementary increase in the number of pixels can as well be performed by upsampling one of the signals and applying low-pass interpolation, but it does not result in an increase of resolving power. It can reduce noise, however. Figure 3 shows the mean-squared error (MSE) of the reconstruction versus the number of images used. The performance increases rapidly with the first six images, but the improvement is marginal beyond that.

In the rest of this paper, we will use four images as input to the super-resolution algorithms. Assuming the low-resolution images were subsampled by almost two, this is the theoretical limit for which our algorithm should be able to reconstruct an image of almost double resolution. In other words, four images are a minimum to have a well-determined system when upsampling by two. Thus, we do not consider the improvement to SNR that the use of more images would bring.

5. RESULTS

The super-resolution algorithm described above is tested in simulations and in practical experiments. A simulation gives complete control over the setup and gives exact knowledge of the registration parameters. It enables us to test the performance of the registration and the reconstruction algorithms separately. The three images that were used in the simulations are shown in Figure 4. In the practical experiment, we tested our algorithm on sets of pictures taken with real digital cameras.




FIGURE 3: MSE of the reconstructed image as a function of the number of images used in the super-resolution algorithm. Six images form a good trade-off between performance and computational complexity.

In both simulation and experiment, we compared our registration algorithm to other registration methods. First, our registration method is compared to the frequency domain algorithms by Marcel et al. [18] and by Lucchese and Cortelazzo [22]. Because Lucchese and Cortelazzo use the same phase correlation method as Marcel et al., the method by Lucchese and Cortelazzo is not included in the simulations where only shifts are used. Next, we also compared it to the spatial domain method based on Taylor expansions by Keren et al. [5]. In the simulations using only shifts, our registration method was also compared to the algorithm by Bergen et al. [24], as it was implemented in the super-resolution imaging software by Farsiu et al. [28]. This was only done for the case of horizontal and vertical shifts, because image rotations are not (yet) implemented in this software.

5.1. Simulation

In the simulation, we started from a high-resolution image, which was considered as the equivalent for continuous space (Figure 5(a)). This image was then multiplied by a Tukey window (Figure 5(b)) to make the image circularly symmetric and thus avoiding all boundary effects. Next, three shifted and rotated copies are created from this high-resolution image. Gaussian zero-mean random variables are used for the shift (pixels) and rotation (degrees) parameters. For the shifts, a standard deviation of 2 is used, while the rotation angles have a standard deviation of 1. The different images are then low-pass filtered using an ideal low-pass filter with cutoff frequency $0.12u_s$ (with u_s the sampling frequency of the high-resolution image) to achieve the setup specified in Section 3.3 and Figure 2. The first of these images (not-moved reference image) will be the reconstruction target for the super-resolution algorithm (Figure 5(c)). And finally, the four images are downsampled by a factor eight.

This results in four low-resolution, shifted and rotated images that can be used as input for the super-resolution algorithm (Figure 5(d)). They are aliasing-free in the frequency band $(-0.04u_s, 0.04u_s)$, and are aliased in the rest of the spectrum as discussed in Section 3.3 and Figure 2. By construction, all shifts are multiples of 0.125, but this information is not used in any of the registration algorithms to keep them generally applicable.

The results using the different algorithms are summarized in Table 1. The registration results with our algorithm are much better than the other frequency domain algorithms by Marcel et al. and Lucchese and Cortelazzo. The motion estimates using the algorithm by Lucchese and Cortelazzo are still accurate up to subpixel precision, while the algorithm by Marcel et al. performs much worse in estimating the rotation angle. Because of this erroneous rotation cancellation, the following motion estimation also fails. The results obtained with the algorithm by Keren et al. are similar to those with our algorithm in both shift and rotation estimation.

Another simulation was also made with only horizontal and vertical shifts. The results of this simulation are listed in Table 2. Our algorithm outperformed the other methods and computed the parameters up to the working precision of the computations. The algorithm by Marcel et al. has clearly lower precision than the other algorithms. The spatial domain algorithms by Keren et al. and by Bergen et al. (as implemented by Farsiu et al.) outperform the frequency domain algorithm by Marcel et al., but have lower precision than our algorithm.

In order to find the same motion parameters in the registration as the parameters that were used to create the images, we need to reverse the order in the registration. In other words, because we first shifted the images and then rotated them in the simulation setup, we need to undo the rotation first and then the shifts. Otherwise, a conversion would have to be made before comparing the two.

5.2. Practical experiment

The different algorithms are also compared in two practical experiments with real images. First, a Leica DC250 black and white digital camera is used, with a Nikon 85 mm optical system. As can be seen from its spatial frequency response [29] (Figure 6), aliasing artifacts can occur with this camera. The camera was firmly fixed on a stable tripod that allows only horizontal and vertical shifts and planar rotations parallel to the image plane.

With this camera setup, four shifted and rotated images of a planar scene are captured (Figures 7(a) and 7(b)¹). The planar scene is a resolution test chart in a plane parallel to the image plane of the camera. These images are then registered using the different registration algorithms to be compared (see Table 3), and a high-resolution image is reconstructed using bicubic interpolation (Figure 8).

¹ When this paper is displayed on a screen or printed, it is possible that additional aliasing is present in the images due to resizing. The full size images are available online [30].




FIGURE 4: High-resolution images used in the simulations. (a) Building, (b) castle, and (c) leaves.




FIGURE 5: Simulation setup. (a) Original high-resolution image. (b) Original image multiplied by a window to make it circularly symmetric. (c) Low-pass filtered image to satisfy the reconstruction conditions. This image is used as reconstruction target. (d) Low-resolution image used as input to the super-resolution algorithm. (e) Reconstructed high-resolution image.

In a second experiment, a set of four color images was taken using a Sigma SD10 digital camera. This camera uses a Foveon X3 sensor, which has three photodetectors (for red, green, and blue) at every pixel location. The camera was held manually in approximately the same position while taking

the pictures, which caused small shifts and rotations between the images (Figure 7). Aliasing is present in the high-frequency regions of the images. The different registration algorithms are then applied to these images and a high-resolution image is reconstructed (see Table 3 and Figure 9).

TABLE 1: Comparison of the average absolute error (μ) and the standard deviation of the error (σ) for the shift and rotation parameters in the different algorithms. 150 simulations were performed for each of the images (Figure 4).

Parameters	Our algorithm		Marcel et al.		Lucchese et al.		Keren et al.	
	μ	σ	μ	σ	μ	σ	μ	σ
Shift (pixels)	0.029	0.038	1.999	11.522	0.327	0.417	0.019	0.027
Rotation angle (deg)	0.126	0.191	19.003	79.086	0.142	0.181	0.053	0.071

TABLE 2: Comparison of the average absolute error (μ) and the standard deviation of the error (σ) for the shift parameter in the different algorithms. 150 simulations with only horizontal and vertical shifts (no rotations) were performed for each of the images.

Parameters	Our algorithm		Marcel et al.		Keren et al.		Bergen et al.	
	μ	σ	μ	σ	μ	σ	μ	σ
Shift (pixels)	3.2e-15	3.9e-15	0.3126	0.3803	4.1e-3	6.0e-3	5.4e-3	7.9e-3




FIGURE 6: Spatial frequency response. (a) Horizontal (dashed line) and vertical (solid line) relative spatial frequency response of the Leica DC250 digital camera used in the experiment. A relative spatial frequency of 1 corresponds to 1017 line widths/picture height and 74 cycles/mm on the image sensor. (b) Relative spatial frequency response (solid line) and its aliased versions (dashed line) after sampling. Although there is no aliasing-free part, the signal-to-aliasing ratio is relatively high for low frequencies, and our algorithm still works.

In these experiments, the estimated cutoff frequency in our algorithm is set as high as possible to obtain reliable results. The shifts are estimated from the central 5% of the frequency domain image, while for the rotation estimation, a disc with $\rho = 0.6$ is used. The use of more information for the rotation estimation than for the shift estimation is required to get sufficient precision (see also Section 6). It can be justified by the fact that in this area, the aliasing component of the sampled signal is smaller than the base spectrum component.

The estimates using the algorithm by Marcel et al. and the shift estimation of the algorithm by Lucchese and Cortelazzo have lower precision due to the size of the images. Because the original images are already relatively large, the required additional upsampling and interpolation require too much memory to be performed on a regular computer. The upsampling is therefore omitted in this case and the shifts are only computed up to pixel level.

Because the exact motion parameters are unknown, it is only possible to compare visually the different reconstructed images. From Figures 8 and 9, it can be seen that with our algorithm, the registration was very accurate. Most aliasing has been removed in both reconstructed images. In the images obtained with the algorithms by Marcel et al. and by Lucchese and Cortelazzo, at least one of the images was badly aligned. Therefore, the reconstructions are also less precise. The results with the algorithm by Keren et al. are comparable to the results with our algorithm.

6. DISCUSSION

A very precise registration algorithm is required for any super-resolution algorithm to work. From the comparison in simulations and with real image sequences, it is clear that our frequency domain algorithm and the spatial domain algorithm by Keren et al. are accurate enough to improve

TABLE 3: Registration parameters for the practical experiments using our algorithm, nonlinear minimization, and the algorithm by Keren et al. Experiment 1 is the experiment with the resolution chart using the Leica camera. Experiment 2 is the experiment with the outdoor scene using the Sigma camera.

Im. pairs	Our algorithm			Marcel et al.			Lucchese et al.			Keren et al.		
	Δx	Δy	θ	Δx	Δy	θ	Δx	Δy	θ	Δx	Δy	θ
Exp. 1												
Im2-Im1	9.24	-3.84	0.9	11.2	-4.20	1.06	9.00	-0.50	1.21	9.27	-3.86	0.92
Im3-Im1	9.74	-2.21	1.2	12.4	0.20	1.39	10.00	2.00	1.68	9.86	-2.29	1.14
Im4-Im1	10.32	-5.00	1.2	12.4	-5.00	1.39	11.25	-0.75	1.63	10.37	-5.06	1.17
Exp. 2												
Im2-Im1	-12.75	-10.34	-0.1	-17	-10	0	-15	-4	-0.53	-12.51	-10.43	-0.01
Im3-Im1	14.65	12.96	0	22	13	0	11	36	1.66	15.22	13.08	0.01
Im4-Im1	-12.08	1.54	-0.1	-18	2	0	-14	-6	-0.63	-12.76	1.74	-0.09




FIGURE 7: Aliased images taken with a real digital camera and used in the practical experiments. (a) One of the four images of the resolution chart taken with the Leica digital camera, and (b) a detail showing the aliasing. (c) One of the four images of a real-life scene taken with the Sigma digital camera, and (d) a detail showing the aliasing. The four images for both experiments are available online [30].

resolution and remove aliasing artifacts (see Figures 8 and 9). Our algorithm performs better than the algorithm by Keren et al. if there is some strong directionality present in the images. The other frequency domain algorithms by Marcel et al. and Lucchese and Cortelazzo perform worse both in the simulations and in the practical experiment.

We can also observe that a bad image registration is fatal for the reconstruction. In such cases, it would be better to reconstruct a larger image from only one of the low-resolution images using interpolation, even though this does not increase the resolution. The artifacts due to bad motion estimation are visually very noticeable.

Our algorithm works best on images with strong frequency content in certain directions (Figures 10(a) and 10(b)). In that case, our algorithm outperforms all other algorithms including the spatial domain algorithm by Keren et al. The accuracy of our rotation estimation (and consequently also of the shift estimation) depends on the presence of some strong directionality in the images. This can be observed in Table 4, where the results from Table 1 for our algorithm are displayed per image. If such frequency directions are not present (Figures 10(c) and 10(d)), the registration performance decreases. The results with our algorithm are then slightly worse than with the algorithm by Keren et al., but still much better than those using the other frequency domain algorithms. This dependence on directionality is related to the projection along radial lines in our rotation estimation algorithm. This highly reduces the computational complexity of the algorithm, as only a one-dimensional correlation is required instead of the regular two-dimensional correlations. However, because of the projection, it is also more subject to errors if there are no strong directions in the image.

Next to the presence of directional frequency content, the size of the low-resolution images also constrains the precision of our rotation estimation algorithm. As the frequency values have to be averaged over a small angle (typically a few degrees), the number of values to be averaged will be very limited for small images. This number of values also varies for different angles (e.g., more values around 0 and 90 degrees, less in between), which biases the computed functions. In Table 5, simulation results with our algorithm are compared for different image sizes. This explains also why we consider a large disc for the rotation estimation, as estimates based on the aliasing-free part alone are not accurate enough.

The super-resolution technique described in Sections 3 and 4 can be applied in many different applications, such as surveillance, consumer digital cameras, aerial photography, and so forth. However, an important limitation to its direct application can be found in the current camera design. Because aliasing is visually so disturbing, most digital camera manufacturers design the optical system of their cameras to




FIGURE 8: Results of the different super-resolution algorithms on the real images of the resolution chart. Zoomed images of the central part are displayed to show the differences better. (a) Our algorithm. (b) Registration algorithm by Marcel et al. (c) Registration algorithm by Lucchese and Cortelazzo. (d) Registration algorithm by Keren et al.




FIGURE 9: Results of the different super-resolution algorithms on the images of the real-life scene. Zoomed images of the central part are displayed to show the differences better. (a) Our algorithm. (b) Registration algorithm by Marcel et al. (c) Registration algorithm by Lucchese and Cortelazzo. (d) Registration algorithm by Keren et al.

remove aliasing. An optical low-pass filter is applied to the image before it is captured to ensure that aliasing cannot occur. Our technique totally relies on the presence of aliasing in

the captured images, so if the images are free of aliasing, our algorithm cannot perform better than a regular interpolation from a single image.




FIGURE 10: Our algorithm works best on images with strong frequency content in a number of directions ((a) and its Fourier transform (b)). If the energy is homogeneously spread among all possible directions (as can be seen in (c) and its Fourier transform (d)), the performance of the motion estimation algorithm decreases.

TABLE 4: Comparison of the average absolute error (μ) and the standard deviation of the error (σ) for the shift and rotation parameters on different images. 150 simulations were performed for each of the images (Figure 4).

Parameters	Image 4(a)		Image 4(b)		Image 4(c)	
	μ	σ	μ	σ	μ	σ
Shift (pixels)	0.041	0.050	0.025	0.027	0.020	0.031
Rotation angle (deg)	0.050	0.034	0.063	0.066	0.265	0.316

The advantage of a frequency domain approach is that the aliasing terms are clearly distinguishable. This makes it much easier to take aliasing into account and accurately estimate the motion, even in the (partially) aliased case. The fact that only low-frequency information was used makes the algorithm also more robust to noise, which is strongest in high frequencies.

The disadvantage of a frequency domain motion estimation algorithm is that it can compute only one set of motion parameters for the whole image. This motion model does not include a scenario when one object in a scene moves and the rest of the scene stays constant. For such cases, a local algorithm is needed.

7. CONCLUSIONS

We presented a new frequency domain method for the registration of a set of low-resolution, aliased images that

TABLE 5: Comparison of the average absolute error (μ) and the standard deviation of the error (σ) for the shift and rotation parameters for different image sizes using our algorithm. 50 simulations were performed for each of the three images (Figure 4).

Input image size	221 \times 221 pixels		884 \times 884 pixels	
	μ	σ	μ	σ
Shift (pixels)	0.0748	0.1010	0.0288	0.0376
Rotation angle (deg)	0.3931	0.5550	0.1261	0.1910

outperforms previous frequency domain registration methods. Planar rotation and translation parameters are precisely estimated based on the low-frequency, aliasing-free part of the images. This image registration technique is then applied to super-resolution imaging to reconstruct a double-resolution image (in each dimension) from a set of aliased images. After the image alignment, bicubic interpolation was used to reconstruct the high-resolution image. This algorithm was compared to some other frequency and spatial domain methods in simulations and practical experiments. Both proved the validity and high precision of our algorithm. If the low-resolution images are sufficiently large and have directionality, our algorithm outperforms the other algorithms and accurate aliasing-free high-resolution images can be reconstructed.

ACKNOWLEDGMENTS

The work presented in this paper was supported in part by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under Grant no. 5005-67322. All the results and figures displayed in this paper are reproducible [31] using the data and Matlab code available online [30]. We thank Urs Schmid from Leica Microsystems AG and Rudy Guttosch from Foveon, Inc. for providing digital cameras, and Sina Farsiu and Peyman Milanfar from UC Santa Cruz for allowing us to use their super-resolution software. We would also like to thank the reviewers for their valuable comments to improve the quality of this paper.

REFERENCES

- [1] R. Y. Tsai and T. S. Huang, "Multiframe image restoration and registration," in *Advances in Computer Vision and Image Processing*, vol. 1, chapter 7, pp. 317–339, JAI Press, Greenwich, Conn, USA, 1984.
- [2] P. Vandewalle, S. E. Süsstrunk, and M. Vetterli, "Super-resolution images reconstructed from aliased images," in *Proceedings of SPIE/IS&T Visual Communications and Image Processing Conference*, T. Ebrahimi and T. Sikora, Eds., vol. 5150 of *Proceedings of SPIE*, pp. 1398–1405, Lugano, Switzerland, 2003.
- [3] P. Vandewalle, S. E. Süsstrunk, and M. Vetterli, "Double resolution from a set of aliased images," in *Proceedings of SPIE/IS&T Electronic Imaging 2004: Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications V*, vol. 5301 of *Proceedings of SPIE*, pp. 374–382, San Jose, Calif, USA, January 2004.

- [4] D. Capel and A. Zisserman, "Computer vision applied to super-resolution," *IEEE Signal Processing Magazine*, vol. 20, no. 3, pp. 75–86, 2003.
- [5] D. Keren, S. Peleg, and R. Brada, "Image sequence enhancement using sub-pixel displacements," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '88)*, pp. 742–746, Ann Arbor, Mich, USA, June 1988.
- [6] 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*, vol. 9, no. 1, pp. 38–50, 1998.
- [7] M. Irani and S. Peleg, "Improving resolution by image registration," *CVGIP: Graphical Models and Image Processing*, vol. 53, no. 3, pp. 231–239, 1991.
- [8] D. Rajan, S. Chaudhuri, and M. V. Joshi, "Multi-objective super-resolution: concepts and examples," *IEEE Signal Processing Magazine*, vol. 20, no. 3, pp. 49–61, 2003.
- [9] M. V. Joshi, S. Chaudhuri, and R. Panuganti, "Super-resolution imaging: use of zoom as a cue," *Image and Vision Computing*, vol. 22, no. 14, pp. 1185–1196, 2004.
- [10] A. J. Patti, M. I. Sezan, and A. Murat Tekalp, "Super-resolution video reconstruction with arbitrary sampling lattices and nonzero aperture time," *IEEE Transactions on Image Processing*, vol. 6, no. 8, pp. 1064–1076, 1997.
- [11] A. Zomet, A. Rav-Acha, and S. Peleg, "Robust super-resolution," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '01)*, vol. 1, pp. 645–650, Kauai, Hawaii, USA, December 2001.
- [12] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, "Fast and robust multiframe super-resolution," *IEEE Transactions on Image Processing*, vol. 13, no. 10, pp. 1327–1344, 2004.
- [13] M. Elad and A. Feuer, "Restoration of a single super-resolution image from several blurred, noisy, and undersampled measured images," *IEEE Transactions on Image Processing*, vol. 6, no. 12, pp. 1646–1658, 1997.
- [14] S. Borman and R. L. Stevenson, "Spatial resolution enhancement of low-resolution image sequences—a comprehensive review with directions for future research," Tech. Rep., Laboratory for Image and Signal Analysis (LISA), University of Notre Dame, Notre Dame, Ind, USA, 1998. Online available: <http://www.nd.edu/~sborman/publications/>.
- [15] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," *IEEE Signal Processing Magazine*, vol. 20, no. 3, pp. 21–36, 2003.
- [16] B. Zitová and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing*, vol. 21, no. 11, pp. 977–1000, 2003.
- [17] B. S. Reddy and B. N. Chatterji, "An FFT-based technique for translation, rotation, and scale-invariant image registration," *IEEE Transactions on Image Processing*, vol. 5, no. 8, pp. 1266–1271, 1996.
- [18] B. Marcel, M. Briot, and R. Murrieta, "Calcul de translation et rotation par la transformation de Fourier," *Traitement du Signal*, vol. 14, no. 2, pp. 135–149, 1997.
- [19] S. P. Kim and W.-Y. Su, "Subpixel accuracy image registration by spectrum cancellation," in *Proceedings of IEEE International Conference Acoustics, Speech, Signal Processing (ICASSP '93)*, vol. 5, pp. 153–156, Minneapolis, Minn, USA, April 1993.
- [20] 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*, vol. 39, no. 10, pp. 2235–2243, 2001.
- [21] H. Foroosh, J. B. Zerubia, and M. Berthod, "Extension of phase correlation to subpixel registration," *IEEE Transactions on Image Processing*, vol. 11, no. 3, pp. 188–200, 2002.
- [22] L. Lucchese and G. M. Cortelazzo, "A noise-robust frequency domain technique for estimating planar roto-translations," *IEEE Transactions on Signal Processing*, vol. 48, no. 6, pp. 1769–1786, 2000.
- [23] 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*, vol. 24, no. 6, pp. 381–395, 1981.
- [24] J. R. Bergen, P. Anandan, K. J. Hanna, and R. Hingorani, "Hierarchical model-based motion estimation," in *Proceedings of 2nd European Conference on Computer Vision (ECCV '92)*, Lecture Notes in Computer Science, pp. 237–252, Santa Margherita Ligure, Italy, May 1992.
- [25] M. Irani, B. Rousso, and S. Peleg, "Computing occluding and transparent motions," *International Journal of Computer Vision*, vol. 12, no. 1, pp. 5–16, 1994.
- [26] J. Gluckman, "Gradient field distributions for the registration of images," in *Proceedings of IEEE International Conference on Image Processing (ICIP '03)*, vol. 3, pp. 691–694, Barcelona, Spain, September 2003.
- [27] A. Papoulis, "Generalized sampling expansion," *IEEE Transactions on Circuits Systems*, vol. 24, no. 11, pp. 652–654, 1977.
- [28] S. Farsiu, M. D. Robinson, and P. Milanfar, "MDSP resolution enhancement software," 2004, Online available: <http://www.soe.ucsc.edu/~milanfar/SR-Software.htm>.
- [29] International Organization for Standardization, "ISO 12233:2000—Photography—Electronic still picture cameras—Resolution measurements," 2000.
- [30] http://lcavwww.epfl.ch/reproducible_research/VandewalleSV05/.
- [31] M. Schwab, M. Karrenbach, and J. Claerbout, "Making scientific computations reproducible," *Computing in Science & Engineering*, vol. 2, no. 6, pp. 61–67, 2000.

Patrick Vandewalle received the M.S. degree in electrical engineering from Katholieke Universiteit Leuven, Belgium in 2001. From 2001 to 2002, he worked as a Research Assistant in the Medical Imaging Lab at the Department of Electrical Engineering (ESAT), KU Leuven. He is currently pursuing the Ph.D. degree in computer, communication, and information sciences at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. His research interests are in signal and image processing, sampling, and digital photography.



Sabine Süstrunk received the B.S. degree in scientific photography from the Swiss Federal Institute of Technology, Zurich, Switzerland, and the M.S. degree in graphic arts publishing, concentration in electronic publishing, from the Rochester Institute of Technology, Rochester, NY, and the Ph.D. in computer science from the University of East Anglia, Norwich, UK in 1987, 1993, and 2005, respectively. She is an Adjunct Professor with the Audiovisual Communications Laboratory at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne,



Switzerland. Her research interests are in digital photography, color image processing, and image quality.

Martin Vetterli received his Engineering degree from ETH Zurich, Switzerland, his M.S. degree from Stanford, and his Ph.D. degree from Ecole Polytechnique Fédérale de Lausanne. In 1986, he joined Columbia University in New York, first with the Center for Telecommunications Research and then with the Department of Electrical Engineering where he was an Associate Professor of electrical engineering. In 1993, he joined the University of California at Berkeley, where he was a Full Professor until 1997. Since 1995, he is a Professor at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, where he headed the Communication Systems Division (1996/1997) and heads the Audiovisual Communications Laboratory. From 2001 to 2004, he directed the National Competence Center in research on mobile information and communication systems. He is also a Vice-President for International Affairs at EPFL since October 2004. He has held visiting positions at ETHZ (1990) and Stanford (1998). His research interests include sampling, wavelets, multirate signal processing, computational complexity, signal processing for communications, digital video processing, and joint source/channel coding. His work won him several prizes (Best Paper Awards from EURASIP in 1984 and of the IEEE Signal Processing Society in 1991 and 1997) and he is the coauthor with J. Kovacevic of the book *Wavelets and Subband Coding* (Prentice-Hall, 1994).

