Available online at www.sciencedirect.com
SCIENCE DIRECT
Pattern Recognition Letters 27 (2006) 287–293
Pattern Recognition Letters
www.elsevier.com/locate/patrec

Resolution enhancement via probabilistic deconvolution of multiple degraded images ☆

F. Šroubek *, J. Flusser

Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Pod vodárenskou věží 4, 182 08, Prague 8, Czech Republic

Available online 29 September 2005

Abstract

We present a maximum a posteriori solution to the problem of obtaining a high-resolution image from a set of degraded low-resolution images of the same scene. The proposed algorithm has the advantage that no prior knowledge of blurring functions is required and it can handle unknown misregistrations between the input images. An efficient implementation scheme of alternating minimizations is presented together with experiments that demonstrate the performance of the algorithm.

Keywords: Image fusion; Multichannel blind deconvolution; Super-resolution; Image restoration; Image registration; MAP estimator

1. Introduction

Image fusion is one of quickly developing advanced methods used in processing of remotely sensed images. The term *fusion* means in general an approach to extraction of information spontaneously adopted in several domains. The goal of image fusion is to integrate complementary multisensor, multitemporal and/or multiview information into one new image containing information the quality of which cannot be achieved otherwise. The term ''quality'' depends on the application requirements.

In remote sensing, image fusion has been mainly used to achieve high spatial and spectral resolutions by combining images from two sensors, one of which has high spatial resolution and the other one high spectral resolution. Typically, a multispectral satellite image (like SPOT XMS for instance) is fused with a high-resolution panchromatic image (SPOT Panchro) or with an aerial image. Most fusion methods are based on image decomposition into low-pass and high-pass bands and on combining the low-pass band of the multispectral image with the high-pass band(s) of the panchromatic image (Duport et al., 1996; Núñez et al., 1999; Scheunders and Backer, 2001; Li et al., 2002; Ranchin et al., 2003). Similar effect can be achieved by transforming the multispectral image into IHS coordinates and replacing intensity component by the panchromatic image (Carper et al., 1990; Chavez et al., 1991).

In this paper, we consider a different formulation of the problem. Assuming two or more low-resolution images from the same sensor (or from different sensors of the same type), our goal is to obtain fused image of higher spatial resolution than the resolution of the input channels. Contrary to the previous case, this task is more complicated because we do not have the high-resolution information in any form. This problem appears in remote sensing very often. Due to the physical limitations of the sensors (see, for instance, Reichenbach et al., 1995 for detailed explanation) and imperfect observational conditions, the acquired images represent only degraded versions of the original scene, where mainly the high-frequency information is

☆ This work was supported by the Grant Agency of the Czech Republic under the project No. 102/04/0155 and partially also by the Czech Ministry of Education under the project No. 1M6798555601 (Research Center DAR).
* Corresponding author. Fax: +420 284680730.
E-mail addresses: sroubekf@utia.cas.cz (F. Šroubek), flusser@utia.cas.cz (J. Flusser).

0167-8655/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.patrec.2005.08.010

suppressed, degraded or missing. Fusion of the low-resolution images is an effective means of breaking the sensor limits and of removing the degradation introduced by atmospheric turbulence, sensor motion, and other factors.

We should point out that this problem appears also outside the area of remote sensing (it has attracted the attention of producers of low-resolution cameras and videos, among others) and has lead to two developing techniques that evolve in parallel: multichannel blind deconvolution (MBD) and super-resolution imaging (SR). First blind deconvolution attempts were based on singlechannel formulations, such as in (Lagendijk and Biedmond, 1991; Reeves and Mersereau, 1992; Chan and Wong, 1998; Haindl, 2000). For a basic survey one can refer to Kundur and Hatzinakos (1996). The problem is extremely ill-posed in the singlechannel framework and lacks any solution in the truly blind case. These methods do not exploit the potential of the multichannel framework, i.e., the missing information about the original image in one channel can be supplemented by the information in other channels. Research on intrinsically multichannel methods has begun fairly recently; refer to Harikumar and Bresler (1999), Giannakis and Heath (2000), Pai and Bovik (2001), Rav-Acha et al. (2000), Panci et al. (2003) and Šroubek and Flusser (2003) for a survey and other references. Such MBD methods brake the limitations of previous techniques and can recover the blurring functions just from the input channels. As regards the SR imaging, see Park et al. (2003) for a basic survey, proposed methods focus mainly on the accurate identification of sub-pixel shifts and on the formulation of different methodologies to find the SR solution. The blurring functions are assumed to be known or estimated by other techniques that are application dependent. An exception to this is given in (Nguyen et al., 2001; Woods et al., 2003), where they carry out blind deconvolution but consider only parametric models of blurs.

A common weakness of the previous techniques is that they need too much a priori information which is not realistic in practice. For instance, they require the knowledge of shape and size of the blurring function, availability of a high-resolution reference frame, or accurate geometric alignment (registration) of the input channels.

In this paper we present a stochastic fusion method that performs MBD and SR simultaneously. We employ and further develop the MBD theory, which we have introduced in (Šroubek and Flusser, 2005). In Section 2, the optimal MBD solution is defined as a maximum a posteriori (MAP) estimate and prior distributions of the original image and blur functions are built. It is the prior distribution of the blurs that distinguishes our approach from singlechannel MAP techniques such as Conan et al. (1998). An alternating minimization (AM) algorithm and a discussion of the SR extension is saved for the end of the section. The main feature of the new fusion method is that it does not require any knowledge of the blurring functions and the input channels might be mutually shifted by an unknown vector. Allowing only translational between-channel misregistration is not a serious limitation. Larger and more complex geometric distortions can be suppressed (usually just up to a small between-image shift) by a proper registration method (there have been hundreds of them investigated, see Zitová and Flusser, 2003 for a survey). Experiments in Section 3 address the issue of performance under different noise levels and misregistration.

2. MAP analysis

Let us assume that the k-th acquired low-resolution image (channel) z_k can be modeled by blurring the "ideal" image u and shifting the result by an unknown vector $(a_k, b_k) = t_k$, i.e.,

$$z_k(x + a_k, y + b_k) = (u * h_k)(x, y) + n_k(x, y), \tag{1}$$

where h_k is an unknown PSF having a character of a low-pass filter, and n_k denotes noise. We assume additive white Gaussian noise (AWGN). This model is a very realistic description of many low-resolution satellite sensors (Reichenbach et al., 1995). In the discrete domain, this degradation model takes the form

$$\boldsymbol{z}_k = \boldsymbol{T}_k \boldsymbol{H}_k \boldsymbol{u} + \boldsymbol{n}_k, \quad k = 1, \ldots, K,$$

where $\boldsymbol{z}_k$, $\boldsymbol{u}$, and $\boldsymbol{n}_k$ are discrete lexicographically ordered equivalents of image functions z_k, u, and n_k, respectively. $\boldsymbol{T}_k$ is a translation operator shifting an image by t_k pixels, i.e., a linear filter with the delta function at the position t_k. One can readily see that the matrix product $\boldsymbol{T}_k \boldsymbol{H}_k = \boldsymbol{G}_k$ defines convolution with a mask $\boldsymbol{g}_k$ that is a shifted version of a mask $\boldsymbol{h}_k$ (discrete representation of h_k). By concatenating the channels, the previous equation can be rewritten in two equivalent forms

$$\boldsymbol{z} = \boldsymbol{G}\boldsymbol{u} + \boldsymbol{n} = \boldsymbol{U}\boldsymbol{g} + \boldsymbol{n}, \tag{2}$$

where $\boldsymbol{z} \equiv [\boldsymbol{z}_1^{\mathrm{T}}, \ldots, \boldsymbol{z}_K^{\mathrm{T}}]^{\mathrm{T}}$, $\boldsymbol{G} \equiv [\boldsymbol{G}_1^{\mathrm{T}}, \ldots, \boldsymbol{G}_K^{\mathrm{T}}]^{\mathrm{T}}$, $\boldsymbol{n} \equiv [\boldsymbol{n}_1^{\mathrm{T}}, \ldots, \boldsymbol{n}_K^{\mathrm{T}}]^{\mathrm{T}}$, $\boldsymbol{g} \equiv [\boldsymbol{g}_1^{\mathrm{T}}, \ldots, \boldsymbol{g}_K^{\mathrm{T}}]^{\mathrm{T}}$, and $\boldsymbol{U}$ is a block-diagonal matrix with K blocks each performing convolution with the image $\boldsymbol{u}$.

Adopting a stochastic approach, the problem of image fusion can be formulated as a MAP estimation. We assume that the images $\boldsymbol{u}$, $\boldsymbol{g}$ and $\boldsymbol{z}$ are random vector fields with given probability density functions (pdf) $p(\boldsymbol{u})$, $p(\boldsymbol{g})$ and $p(\boldsymbol{z})$, respectively, and we look for such realizations of $\boldsymbol{u}$ and $\boldsymbol{g}$ which maximize the a posteriori probability $p(\boldsymbol{u}, \boldsymbol{g}|\boldsymbol{z})$. According to the Bayes rule, the relation between a priori probabilities $p(\boldsymbol{u})$, $p(\boldsymbol{g})$ and the a posteriori probability is $p(\boldsymbol{u}, \boldsymbol{g}|\boldsymbol{z}) \propto p(\boldsymbol{z}|\boldsymbol{u}, \boldsymbol{g}) p(\boldsymbol{u}) p(\boldsymbol{g})$. The conditional pdf $p(\boldsymbol{z}|\boldsymbol{u}, \boldsymbol{g})$ follows from (2) and from our assumption of AWGN, i.e.,

$$p(\boldsymbol{z}|\boldsymbol{u}, \boldsymbol{g}) \propto \exp\left\{-\frac{1}{2}(\boldsymbol{z} - \boldsymbol{G}\boldsymbol{u})^{\mathrm{T}} \Sigma^{-1} (\boldsymbol{z} - \boldsymbol{G}\boldsymbol{u})\right\},$$

where Σ is the noise diagonal covariance matrix with $\{\sigma_k^2\}_{k=1}^K$ on the corresponding positions on the main diagonal. If the same noise variance σ^2 is assumed in each channel, Σ^{-1} reduces to a scalar σ^{-2}.

2.1. A priori distribution of the original image

A general model for the prior distribution $p(u)$ is a Markov random field which is characterized by its Gibbs distribution given by $p(u) \propto \exp(-F(u)/\lambda)$, where λ is a constant and F is called the *energy function*. One can find various forms of the energy function in the literature, however, the most promising results have been achieved for variational integrals. The energy function then takes the form

$$F(u) = \int \phi(|\nabla u|), \tag{3}$$

where ϕ is strictly convex, nondecreasing function that grows at most linearly. Examples of $\phi(s)$ are s (total variation), $\sqrt{1+s^2}-1$ (hypersurface minimal function) or $\log(\cosh(s))$. Nonconvex functions, such as $\log(1+s^2)$, $s^2/(1+s^2)$ or $\arctan(s^2)$ (Mumford–Shah functional), may behave in an unpredictable manner but since they provide better results for segmentation problems, they are often used as well. The energy function based on the variational integral is highly nonlinear and to overcome this difficulty we follow a half-quadratic scheme described in (Charbonnier et al., 1997) which introduces an auxiliary variable. A special attention must be paid to the discretization of the image gradient ∇u and relaxation of ϕ. In addition, we confine the distribution to an amplitude constraint set $C_u \equiv \{u | \alpha \leqslant u \leqslant \beta\}$ with amplitude bounds derived from the input images, typically $\alpha = 0$ and $\beta = 255$. We thus define the prior distribution as

$$p(\boldsymbol{u}) = \begin{cases} \dfrac{1}{Z} \exp\left\{-\dfrac{1}{2\sigma_u^2}\boldsymbol{u}^{\mathrm{T}}\boldsymbol{L}(v)\boldsymbol{u}\right\} & \text{if } \boldsymbol{u} \in C_u, \\ 0 & \text{otherwise,} \end{cases}$$

where Z is the partition function, σ_u^2 denotes the image variance, $\boldsymbol{u}^{\mathrm{T}}\boldsymbol{L}(v)\boldsymbol{u}$ represents the discretization of (3) and v is the auxiliary variable introduced by the half-quadratic scheme, which is calculated as

$$v(x,y) = \frac{\phi'(|\nabla u(x,y)|)}{|\nabla u(x,y)|}. \tag{4}$$

Matrix $\boldsymbol{L}(v)$ is a positive semidefinite block tridiagonal matrix constructed by v. It performs shift-variant convolution with v.

2.2. A priori distribution of the blurs

The shape of the prior distribution $p(\boldsymbol{g})$ can be derived from the fundamental multichannel constraint stated in (Harikumar and Bresler, 1999; Giannakis and Heath, 2000). Let $\boldsymbol{Z}_k$ denote the convolution matrix with the degraded image $\boldsymbol{z}_k$. If noise $\boldsymbol{n}_k$ is zero and the original channel masks $\{\boldsymbol{h}_k\}$ are *weakly coprime*, i.e., their only common factor is a scalar, then all solutions $\{\hat{\boldsymbol{g}}_k\}$ to

$$\boldsymbol{Z}_i\hat{\boldsymbol{g}}_j - \boldsymbol{Z}_j\hat{\boldsymbol{g}}_i = \boldsymbol{0}, \quad 1 \leqslant i < j \leqslant K \tag{5}$$

have the following forms. Let S_g denote the sum of the maximum blur size and the maximum shift between the channels. If S_g is known the solution equals $\{\alpha\boldsymbol{g}_k\}$ for any scalar α. If S_g is not known, it must be first estimated and two distinct situations arise. If S_g is underestimated, zero vector is the only solution of (5). If S_g is overestimated, then the space of all solutions of (5) contains the correct masks $\{\boldsymbol{g}_k\}$ and the dimensionality of this solution space is proportional to the degree of the overestimation. Further stacking the system of Eq. (5), we obtain

$$\mathscr{Z}\hat{\boldsymbol{g}} = \boldsymbol{0}, \tag{6}$$

where $\hat{\boldsymbol{g}} \equiv [\hat{\boldsymbol{g}}_1^{\mathrm{T}}, \ldots, \hat{\boldsymbol{g}}_K^{\mathrm{T}}]^{\mathrm{T}}$. If the noise term $\boldsymbol{n}_k$ is present, it follows from (2) that the left-hand side of (6) equals a realization of a Gaussian process of zero mean and covariance $\mathscr{C} = \mathscr{G}\Sigma\mathscr{G}^{\mathrm{T}}$, where $\mathscr{G}$ takes the form of $\mathscr{Z}$ in (6) with $\boldsymbol{Z}_i$ replaced by $\boldsymbol{G}_i$.

It is desirable to include also other prior knowledge about the blurs, such as positivity or constant energy. We therefore define a set of admissible solutions as $C_g \equiv \{\boldsymbol{g} | g_k(x,y) \geqslant 0 \wedge \sum_{x,y} g_k(x,y) = 1,\ k = 1, \ldots, K\}$ and propose the following prior distribution:

$$p(\boldsymbol{g}) = \begin{cases} \dfrac{1}{Z} \exp\left\{-\dfrac{1}{2}\boldsymbol{g}^{\mathrm{T}}\mathscr{Z}^{\mathrm{T}}\mathscr{C}^{-1}\mathscr{Z}\boldsymbol{g}\right\} & \text{if } \boldsymbol{g} \in C_g, \\ 0 & \text{otherwise.} \end{cases}$$

The inverse of the matrix $\mathscr{C}$ is not trivial. In addition, the matrix is constructed by the blurs $\boldsymbol{g}$ that are unknown. To overcome this difficulty, we approximate $\mathscr{C}$ by a diagonal matrix $\mathscr{D}$ such that $\mathrm{diag}(\mathscr{D}) = \mathrm{diag}(\mathscr{C})$, where $\mathrm{diag}(\cdot)$ denotes the main diagonal of the matrix. The elements of $\mathscr{D}$ take the form $\sigma_i^2\|\boldsymbol{g}_j\|^2 + \sigma_j^2\|\boldsymbol{g}_i\|^2$ for $1 \leqslant i < j \leqslant K$. The value of $\|\boldsymbol{g}_i\|^2$ is not known in advance, but a good initial approximation can be given. Since $\boldsymbol{g} \in C_g$,

$$\frac{1}{\prod S_g} \leqslant \|\boldsymbol{g}_i\|^2 \leqslant 1$$

and we use the bottom limit for $\|\boldsymbol{g}_i\|^2$.

2.3. AM-MAP algorithm

The MAP estimation is given by

$$\{\hat{u}, \hat{g}\} = \arg\min_{\boldsymbol{u}\in C_u, \boldsymbol{g}\in C_g} \left\{(\boldsymbol{z} - \boldsymbol{G}\boldsymbol{u})^{\mathrm{T}}\Sigma^{-1}(\boldsymbol{z} - \boldsymbol{G}\boldsymbol{u}) + \frac{1}{\sigma_u^2}\boldsymbol{u}^{\mathrm{T}}\boldsymbol{L}(v)\boldsymbol{u} + \boldsymbol{g}^{\mathrm{T}}\mathscr{Z}^{\mathrm{T}}\mathscr{D}^{-1}\mathscr{Z}\boldsymbol{g}\right\}. \tag{7}$$

Such problems can be solved by means of genetic algorithms, e.g. simulated annealing. In this paper we have adopted an approach of alternating minimizations over $\boldsymbol{u}$ and $\boldsymbol{g}$. The advantage of this scheme lies in its simplicity. Each term in (7) is convex and the derivatives w.r.t. $\boldsymbol{u}$ and $\boldsymbol{g}$ can be easily calculated. The proposed AM-MAP algorithm alternates between two steps

1. $\boldsymbol{u} = \left(\boldsymbol{G}^{\mathrm{T}}\Sigma^{-1}\boldsymbol{G} + \dfrac{1}{\sigma_u^2}\boldsymbol{L}(v)\right)^{-1}\boldsymbol{G}^{\mathrm{T}}\Sigma^{-1}\boldsymbol{z} \wedge \boldsymbol{u} \in C_u,$
2. $\boldsymbol{g} = (\boldsymbol{U}^{\mathrm{T}}\Sigma^{-1}\boldsymbol{U} + \mathscr{Z}^{\mathrm{T}}\mathscr{D}^{-1}\mathscr{Z})^{-1}\boldsymbol{U}^{\mathrm{T}}\Sigma^{-1}\boldsymbol{z} \wedge \boldsymbol{g} \in C_g.$

The algorithm in its essence is similar to our previously proposed one in (Šroubek and Flusser, 2003). In step 1

the flux variable v is updated according to (4). Inversion of the matrix in step 1 cannot be carried out directly because of the matrix size. Instead, we use the method of conjugate gradients and project the solution to C_u. Both $\boldsymbol{U}^{\mathrm{T}}\Sigma^{-1}\boldsymbol{U}$ and $\mathscr{X}^{\mathrm{T}}\mathscr{D}^{-1}\mathscr{X}$ have the size squared proportional to the blur size and can be constructed directly without building the individual matrices $\boldsymbol{U}$ and $\mathscr{X}$ that have much larger size. Then to find $\boldsymbol{g}$, we use a subspace trust region method based on the interior-reflective Newton method (Matlab implements the method in *fmincon* function). This is a constrained minimization problem that confines the solution to C_g and guarantees that assumptions about blurs such as positivity and energy preservation are satisfied.

2.4. Super-resolution extension

So far, we have discussed solely the MBD problem. One can extend the method to perform SR by introducing a matrix $\boldsymbol{D}$ that is called a *decimation operator*. Operator $\boldsymbol{D}$ model the low-resolution acquisition of digital sensors by performing convolution with a uniform mask (spatial averaging) and downsampling of images. In the discrete case, e.g. for SR by factor of 2 in both directions, $\boldsymbol{D}$ returns every second pixel of convolution with a 2×2 uniform mask. Introducing $\boldsymbol{D}$ into (2) the acquisition model becomes:

$$\boldsymbol{z} = \boldsymbol{DGu} + \boldsymbol{n}.$$

The steps in the AM-MAP algorithm are the same except we replace $\boldsymbol{G}$, $\boldsymbol{U}$ and $\mathscr{X}$ with $\boldsymbol{DG}$, $\boldsymbol{DU}$ and $\mathscr{X}\boldsymbol{D}$, respectively. In each term $\boldsymbol{D}$ is constructed so that its size matches the left and right arguments.

One must supply the blur size to the algorithm. An important feature is that an accurate estimation is not necessary; we must only guarantee that the blur size is not underestimated. In addition, the noise covariance Σ and the image variance σ_u^2 are mandatory in the algorithm. However, if noise has the same variance σ_n^2 in every channel, the MAP expression (7) is simplified and only the signal to noise ratio σ_u^2/σ_n^2 is required.

3. Experimental results

Experiments were carried out on two simulated scenarios, seemingly acquired at different time instances. For instance, one can imagine they were captured by a low-resolution sensor similar to AVHRR. In the first simulated scenario, we have analyzed the capability of the algorithm to perform MBD and SR simultaneously in comparison to standard MBD followed by linear interpolation. The second simulated scenario demonstrates robustness of the algorithm to misregistration of the input channels under different noise levels. In all the experiments motion blurs of different extent, width and direction were used to simulate the most common type of degradation in remote sensing. However, it is important to underline that the proposed technique is completely blind and that any type of blur can be considered. The fact that we use the motion blurs was not taken into account in any way. Additionally, the correct noise and image variances were assumed in all the experiments.

3.1. Deconvolution and super-resolution

The data source for the first simulation (playing the role of the "ideal" image) was the 300×300 SPOT HRV image covering the north-western part of Prague (Czech capital); see Fig. 1(a). We chose this image as a typical representation of urban areas. To simulate low-resolution acquisitions, the image was blurred by randomly generated 6×6 motion masks, corrupted by AWGN of SNR = 50 dB and resolution decimated by factor of two to obtain images of size 150×150. Six such images were generated and used as input channels to guarantee a sufficient number of equations. Note that to increase resolution by factor of two, the minimum number of channels is four to have at least as many equations as unknowns. Fig. 1(b) shows one of the low-resolution channels. The result of fusion using the blind deconvolution approach in (Šroubek and Flusser, 2003) and applying linear interpolation afterward is depicted in Fig. 1(c). The proposed algorithm with the SR extension performs better and gives a more accurate representation of the original image, as illustrated in Fig. 1(d). The performance boost, however, diminishes as the noise level increases. We demonstrate this in Fig. 2, which shows the results of the same experiment but for SNR = 10 dB.

3.2. Overestimation of the blur size

To demonstrate that the AM-MAP algorithm performs well even if the size of the blurs is overestimated, we have conducted in the second scenario the following experiments. A 300×300 image of agricultural areas in Fig. 4(c) was degraded with three random motion blurs of size 3×3 and with AWGN of SNR = 10, 30, 50 dB, respectively. The original image was recovered from each image triplet using the algorithm without the SR extension and for four different blur sizes: 3×3, 5×5, 7×7 and 9×9. The percentage mean squared error of the estimated image $\hat{\boldsymbol{u}}$ defined as $\mathrm{PMSE}(\hat{\boldsymbol{u}}) = 100\|\hat{\boldsymbol{u}}-\boldsymbol{u}\|/\|\boldsymbol{u}\|$, was used as the evaluation measure at each iteration. Calculated PMSEs are summarized in Fig. 3. One can see that the convergence rate is virtually not affected by the blur size. Only in the case of SNR = 50 dB, we observe a decrease of the convergence rate for extensive overestimation. This performance drop is not visible for higher levels of noise since it is obscured by the uncertainty implied by noise, and naturally, the error of reconstruction increases with more noise.

Finally, we have conducted an experiment that simulates more accurately conditions in real applications and that utilizes blur size overestimation. The same image in Fig. 4(c) was blurred by two randomly generated 5×5 motion masks and corrupted by AWGN of SNR = 50 dB. The degraded images were mutually shifted by 5 pixels in both directions see Fig. 4(a) and (b) to simulate channel

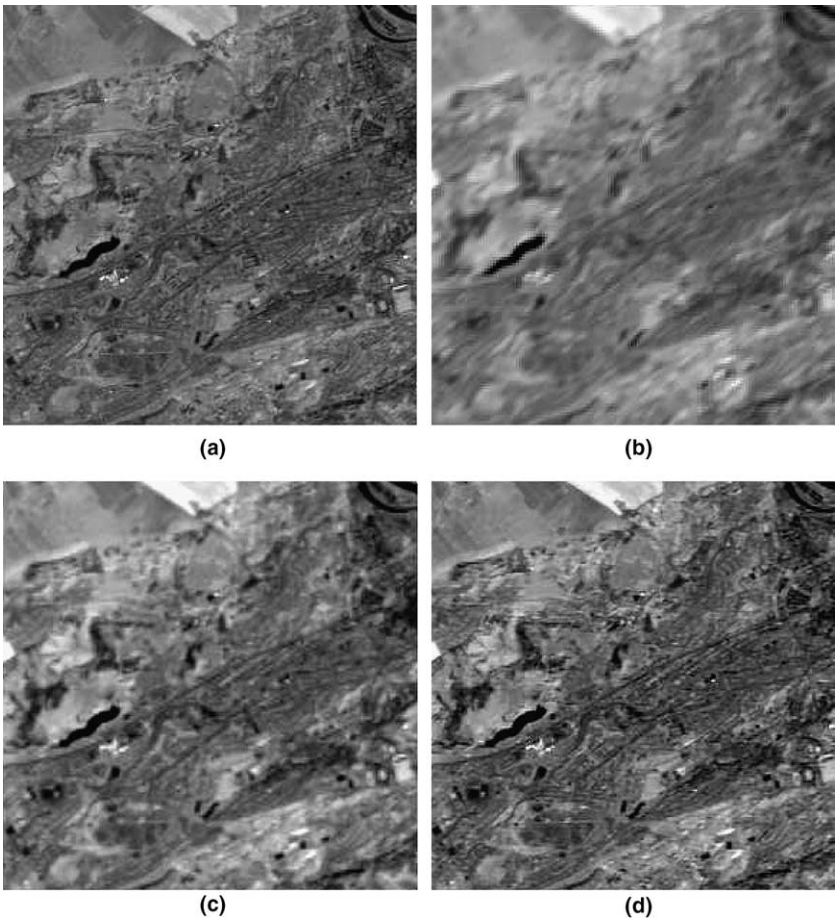

Fig. 1. Fusion of urban-area images without noise: (a) original "ideal" image of size 300×300; (b) simulated blurred and low-resolution images of size 150×150. Six images of type (b), but each with a slightly different blur, were fused using two techniques: (c) result of fusion using the standard MBD followed by linear interpolation; (d) result of fusion using AM-MAP with the SR extension. For the visualization purposes, resolution of image (b) was doubled with the nearest neighbor interpolation.

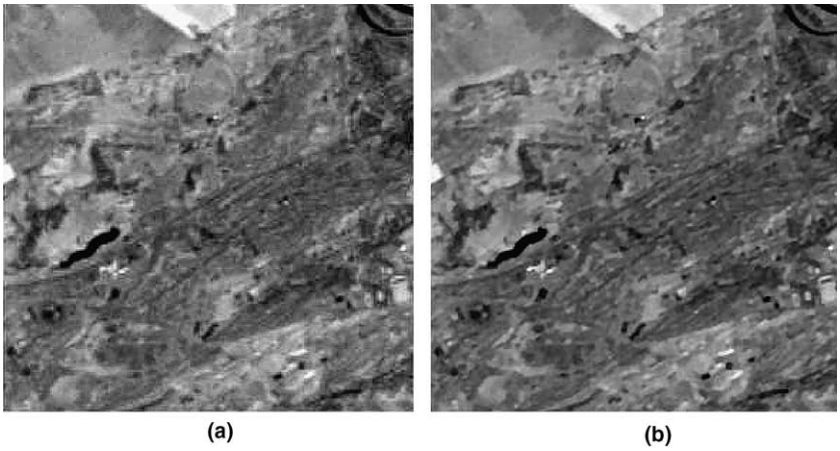

Fig. 2. Fusion of urban-area images with noise: Six images as in Fig. 1(b), but with noise of SNR = 10 dB, were fused using two techniques: (a) result of fusion using the standard MBD followed by linear interpolation; (b) result of fusion using AM-MAP with the SR extension.

misregistration. The AM-MAP algorithm was initialized with the overestimated blur size 12×12 and the result is in Fig. 4(d). The fused image is by visual comparison much sharper than the input channels and is fully comparable to the "ideal" images, which demonstrates an excellent performance of the method.

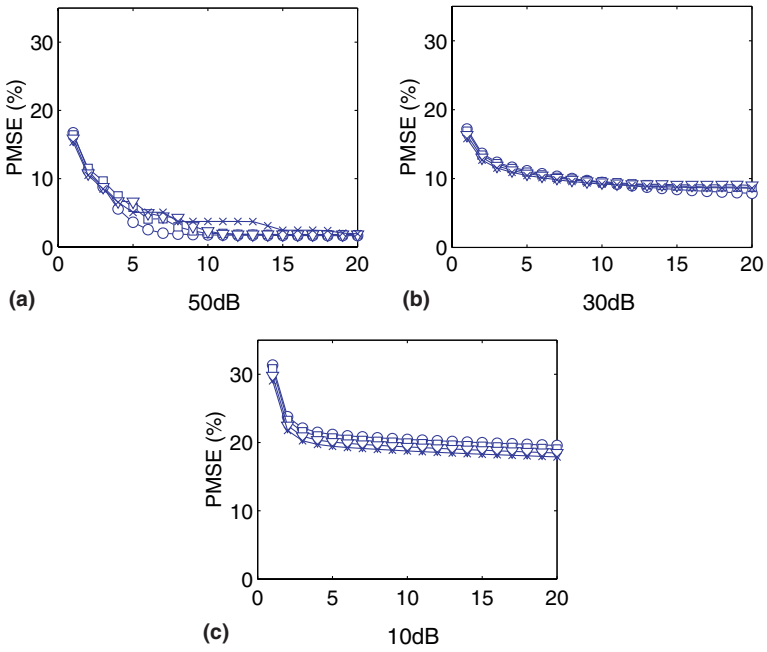

Fig. 3. PMSE of the estimated image as a function of iteration. The image in Fig. 4(c) was degraded with three random blurs of size 3×3 and with AWGN of SNR (a) 50 dB, (b) 30 dB, and (c) 10 dB, respectively. The algorithm was restarted with four different blur sizes: ($\bigcirc$) 3×3 (correct size), ($\square$) 5×5, ($\triangledown$) 7×7, and ($\times$) 9×9.

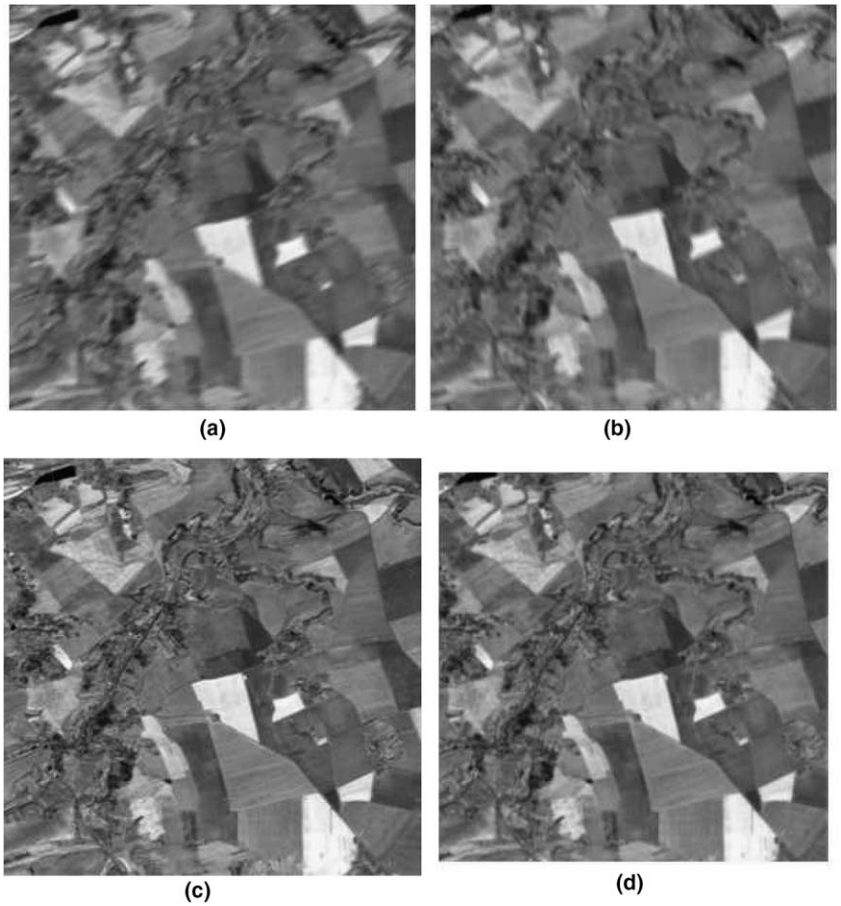

Fig. 4. AM-MAP fusion of agricultural-area images: (a, b) simulated low-resolution images; (c) "ideal" image for comparison; (d) result of fusing (a, b). The fused image is of smaller size since only the overlapping areas of the input images can be fused.

4. Conclusions

We have developed an iterative fusion algorithm that recovers a high-resolution image from misaligned and blurred input channels. The fusion problem is formulated as the MAP estimation with the prior probabilities derived from the variational integral and from the mutual relation of coprime channels. The presented experiments indicate that this approach provides high-quality fused images, fully comparable to the "ideal" ones. Solving the super-resolution and blind deconvolution problems simultaneously is a pioneering step in the field of image restoration.

References

Carper, J., Lillesand, T., Kiefer, R., 1990. The use of Intensity-Hue-Saturation transformations for merging SPOT panchromatic and multispectral image data. Photogramm. Eng. Remote Sens. 56 (4), 459–467.

Chan, T., Wong, C., 1998. Total variation blind deconvolution. IEEE Trans. Image Process. 7 (3), 370–375.

Charbonnier, P., Blanc-Feraud, L., Aubert, G., Barlaud, M., 1997. Deterministic edge-preserving regularization in computed imaging. IEEE Trans. Image Process. 6 (2), 298–311.

Chavez, P., Sides, S., Anderson, J., 1991. Comparison of three different methods to merge multiresolution and multispectral data: Landsat TM and SPOT panchromatic. Photogramm. Eng. Remote Sens. 57, 295–303.

Conan, J., Mugnier, L., Fusco, T., Michau, V., Rousset, G., 1998. Myopic deconvolution of adaptive optics images by use of object and point-spread function power spectra. Appl. Opt. 37 (21), 4614–4622.

Duport, B., Girel, J., Chassery, J., Pautou, G., 1996. The use of multiresolution analysis and wavelets transform for merging SPOT panchromatic and multispectral image data. Photogramm. Eng. Remote Sens. 69 (9), 1057–1066.

Giannakis, G., Heath, R., 2000. Blind identification of multichannel FIR blurs and perfect image restoration. IEEE Trans. Image Process. 9 (11), 1877–1896.

Haindl, M., 2000. Recursive model-based image restoration. Proc. 15th Internat. Conf. on Pattern Recognition, vol. III. IEEE Press, pp. 346–349.

Harikumar, G., Bresler, Y., 1999. Perfect blind restoration of images blurred by multiple filters: Theory and efficient algorithms. IEEE Trans. Image Process. 8 (2), 202–219.

Kundur, D., Hatzinakos, D., 1996. Blind image deconvolution. IEEE Signal Process. Mag. 13 (3), 43–64.

Lagendijk, R.L., Biedmond, J., 1991. Iterative Identification and Restoration of Images. Kluwer, Boston, MA.

Li, S., Kwok, J., Wang, Y., 2002. Using the discrete wavelet frame transform to merge Landsat TM and SPOT panchromatic images. Inform. Fusion 3, 17–23.

Nguyen, N., Milanfar, P., Golub, G., 2001. Efficient generalized cross-validation with applications to parametric image restoration and resolution enhancement. IEEE Trans. Image Process. 10 (9), 1299–1308.

Núñez, J., Otazu, X., Fors, O., 1999. Image fusion with additive multiresolution wavelet decomposition. Applications to SPOT + Landsat images. J. Opt. Soc. Amer. A 16 (3), 467–474.

Pai, H.-T., Bovik, A., 2001. On eigenstructure-based direct multichannel blind image restoration. IEEE Trans. Image Process. 10 (10), 1434–1446.

Panci, G., Campisi, P., Colonnese, S., Scarano, G., 2003. Multichannel blind image deconvolution using the bussgang algorithm: Spatial and multiresolution approaches. IEEE Trans. Image Process. 12 (11), 1324–1337.

Park, S., Park, M., Kang, M., 2003. Super-resolution image reconstruction: A technical overview. IEEE Signal Process. Mag. 20 (3), 21–36.

Ranchin, T., Aiazzi, B., Alparone, L., Baronti, S., Wald, L., 2003. Image fusion—the ARSIS concept and some successful implementation schemes. ISPRS J. Photogramm. Remote Sens. 58, 4–18.

Rav-Acha, A., Peleg, S., 2000. Restoration of multiple images with motion blur in different directions. In: IEEE Workshop on Applications of Computer Vision (WACV), pp. 22–27.

Reeves, S., Mersereau, R., 1992. Blur identification by the method of generalized cross-validation. IEEE Trans. Image Process. 1 (3), 301–311.

Reichenbach, S.E., Koehler, D.E., Strelow, D.W., 1995. Restoration and reconstruction of AVHRR images. IEEE Trans. Geosci. Remote Sens. 33, 997–1007.

Scheunders, P., Backer, S.D., 2001. Fusion and merging of multispectral images with use of multiscale fundamental forms. J. Opt. Soc. Amer. A 18 (10), 2468–2477.

Šroubek, F., Flusser, J., 2003. Multichannel blind iterative image restoration. IEEE Trans. Image Process. 12 (9), 1094–1106.

Šroubek, F., Flusser, J., 2005. Multichannel blind deconvolution of spatially misaligned images. IEEE Trans. Image Process. 14 (7), 874–883.

Woods, N., Galatsanos, N., Katsaggelos, A., 2003. EM-based simultaneous registration, restoration, and interpolation of super-resolved images, In: Proc. IEEE ICIP, vol. 2, pp. 303–306.

Zitová, B., Flusser, J., 2003. Image registration methods: A survey. Image Vision Comput. 21, 977–1000.