Multidim Syst Sign Process (2009) 20:375–384
DOI 10.1007/s11045-008-0072-1

Superresolution reconstruction using nonlinear gradient-based regularization

Xin Zhang · Edmund Y. Lam

Received: 21 May 2008 / Revised: 29 October 2008 / Accepted: 13 November 2008 /
Published online: 13 December 2008

Abstract This paper discusses the problem of superresolution reconstruction. To preserve edges accurately and efficiently in the reconstruction, we propose a nonlinear gradient-based regularization that uses the gradient vector field of a preliminary high resolution image to configure a regularization matrix and compute the regularization parameters. Compared with other existing methods, it not only enhances the spatial resolution of the resulting images, but can also preserve edges and smooth noise to a greater extent. The advantages are shown in simulations and experiments with synthetic and real images.

Keywords Superresolution · Inverse problem · Image reconstruction

1 Introduction

In most applications of electronic imaging, images with high resolution (HR) are desirable and often required. HR means that a large amount of pixels are employed to display an object, and hence edges and details are presented using sufficient pixels, which may be critical in various applications. For example, medical diagnosis and treatment require use of HR images to examine the anatomical and structural information of human bodies and distinguish organs and parts in pathological changes from normal ones at the initial stages. Meanwhile, HR satellite images may be helpful to recognize objects from similar ones. As a processing method, superresolution reconstruction is a possible means to increase the resolution of images obtained from common imaging processes. In recent years, the technique has been applied to various imaging modalities, including optical coherence tomography (OCT) (Blu et al. 2002), digital holography (Mico et al. 2006), MRI (Peled and Yeshurun 2001)

X. Zhang (✉) · E. Y. Lam
Department of Electrical and Electronic Engineering, The University of Hong Kong,
Pokfulam Road, Kowloon, Hong Kong
e-mail: xinzhang@eee.hku.hk

E. Y. Lam
e-mail: elam@eee.hku.hk

and PET (Kennedy et al. 2006). The mathematical technique for superresolution reconstruction in fact originated earlier in the 1980s, as superresolution was applied to distinct types of natural images and videos. Noise involved in the reconstruction has been discussed before (Lam 2003) and many iterative techniques have also been proposed, differing widely in their implementation complexity, memory, and computational speed (Elad and Feuer 1997; Ng et al. 2007; Park et al. 2003; Chan et al. 2007). Nevertheless, the techniques all aimed to efficiently and accurately utilize the a priori knowledge to achieve a good reconstruction. The simplest and most intuitive knowledge of an image is the global smoothness of its spectrum, which gives rise to Tikhonov regularization. Though easy to implement, the resulting image may often not be very sharp, especially when the regularization parameter is set at a high value.

In recent years, many edge-preserving methods have been proposed, such as half-quadratic regularization (HQR) (Charbonnier et al. 1997), directional regularization (Lee et al. 2003) and sparsity regularization in optical tomography (Cao et al. 2007). They are good at recovering edges, but have issues of either heavy computation or specialization. Other than that, adaptive regularization (Vanzella et al. 2004; Watzenig et al. 2004) has also been applied on image restoration and reconstruction. To improve the image spatial resolution while at the same time reducing the computational complexity, in this paper we propose a nonlinear gradient-based regularization, which employs the gradient vector field of a preliminary HR image to configure the regularization matrix and parameter. The regularization is used to form a constrained minimization problem so as to obtain a good solution for the reconstruction. Shin et al. (2001) studied this regularization in noise smoothing with discontinuities and Wang et al. (2005) used it to estimate fingerprint orientation fields, but to our knowledge this has not been applied to superresolution reconstruction.

This paper is structured as follows. In Sect. 2, we describe the model we use in the reconstruction. Two related regularization methods are reviewed for the reconstruction problem. In Sect. 3, the proposed regularization technique and procedure of implementation are presented. Section 4 shows the results of experiments on five images. Section 5 concludes this paper.

2 Observation model and regularization

In image processing applications, a true image f can be related to degraded data y through a linear model of the form

$$y = Hf + n, \tag{1}$$

where H is the system matrix and n denotes noise.

Regularization is a well-established technique for dealing with ill-posed inverse problem in Eq. 1. It imposes the assumption that the true image is reasonably smooth. Thus, a solution can be stabilized by imposing additional smoothness information. We consider the generalized regularization. A penalty term is added to the problem to filter out rough components. The regularized solution f_{reg} takes the form

$$f_{\text{reg}} = \arg\min_f \Big(\|y - Hf\|_R^2 + \alpha J(f) \Big), \tag{2}$$

where $\|y - Hf\|_R^2 = (y - Hf)^T R(y - Hf)$, and R is a diagonal matrix that locally weights the restoration process. This term in Eq. 2 is designed to ensure the fidelity of the solution.

J is a regularization term, i.e., a function representing the method of regularization. It works as constraints and is configured according to the a priori information. Generally, we prefer to set up J as a quadratic form. Then, the minimization is convex and differentiable, with a unique solution. α is the regularization parameter and is set to balance the fidelity and the regularization. The choice of α and configuration of J is critical to the success of a regularization. Among regularization techniques mentioned in Sect. 1, Tikhonov regularization is the simplest to implement, where

$$f_{\text{reg}} = \arg\min_{f} \left(\|y - Hf\|_2^2 + \alpha \|Cf\|_2^2 \right). \tag{3}$$

C is Laplacian operator working as a high-pass filter and $\|\cdot\|_2$ represents L_2 norm.

A variety of regularizations exist, such as total variation regularization and iterative regularization. Here we would like to introduce HQR, because of its significant advantage at preserving edges. HQR was first defined by Geman and Yang (2005). Charbonnier et al. (1997) later proposed to use it in image reconstruction. The regularization introduces an auxiliary variable to solve the minimization problem (2). The variable helps not only mark edges for preservation, but also make the problem half-quadratic, i.e., quadratic with respect to image variable when the auxiliary variable is fixed. Later, Nikolova and Ng (2001) and Deriche et al. (2003) applied this technique to image reconstruction and restoration. The regularized solution is

$$f_{\text{reg}} = \arg\min_{f} \left(\|y - Hf\|_2^2 + \alpha \left[\sum_{i=1}^{4N^2} \varphi(\nabla f_i, s_i) \right] \right). \tag{4}$$

φ is known as the potential function in the Bayesian framework. It determines the regularization imposed on every value of the first-order difference ∇f, which is used to detect the discontinuities of an image. In order to preserve edges, φ is required to satisfy three conditions (Charbonnier et al. 1997)

1. $\dfrac{\varphi'(t)}{2t}$ is continuous and strictly decreasing on $[0, +\infty)$,
2. $\lim\limits_{t \to +\infty} \dfrac{\varphi'(t)}{2t} = 0$,
3. $\lim\limits_{t \to 0^+} \dfrac{\varphi'(t)}{2t} = M$, where $0 < M < +\infty$.

The design of regularization satisfies the conditions of preserving edges in a reconstruction. It works well to recover discontinuities in superresolution reconstruction. The auxiliary variable s ensures edge-preservation in the regularization, but it involves an extra configuration process. The additional process consumes a lot of time to configure every value of vector s. Moreover, a complicated potential function $\varphi(s)$ can add more computation into the discretization and configuration. Thus, HQR usually takes a long time to perform a superresolution reconstruction.

3 Gradient-based regularization

In order to preserve edges efficiently, we propose a nonlinear gradient-based regularization, where the regularization matrix and parameter are calculated locally according to the magnitude and direction of a gradient vector field. Tikhonov regularization demonstrates that a

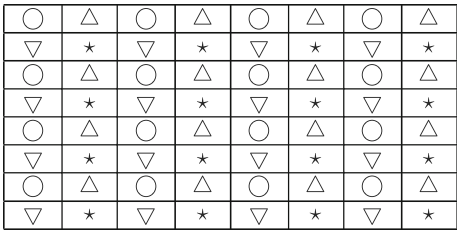

Fig. 1 The diagram of a preliminary HR image. ◯ is an element in the first LR image, △ the second LR, ▽ the third LR image and ⋆ the fourth one

larger value of regularization parameter α leads to a smoother solution. To preserve edges, it is necessary to assign a smaller value to the parameter when it operates on edge pixels. On the other hand, to suppress noise, it is required to set α at a bigger value. The gradient vector field of a preliminary HR image works well to balance the two requirements. The magnitude of the field represents the difference of adjacent pixels. It can sketch the discontinuities of the image and determine whether a pixel is located on an edge or within a region. The direction of the field is utilized to present the orientation of dominant difference around a pixel. After recognizing the edge pixels and their directions, a smaller parameter will be assigned to preserve these edges. Since the configuration of regularization matrix and parameters takes place at the gradient vector field, we can make good use of matrix operations, instead of setting them one by one. So the proposed regularization is also efficient during calculating a reconstruction.

We discuss our algorithm in the most fundamental superresolution reconstruction setting for pedagogical reasons, noting that it can also apply to other configurations with different number and location of the LR images. Assume that four LR images are observed and taken at the same view, but with subpixel spatial shifts. The acquisition process is that the first two images are obtained with no and half pixel shifts in the x direction respectively, and then the other two images are obtained with an additional half pixel shift in the y direction. By reorganizing the pixel values of the four LR images, we can form a preliminary HR image as shown in Fig. 1, in which ◯, △, ▽ and ⋆ represent elements in the four LR images, respectively.

A preliminary image is used to calculate the gradient vector field. Sobel masks perform the 2-D spatial gradient measurement on the preliminary image, and emphasize discontinuities corresponding to edges with a high spatial gradient. Then, according to the direction of the gradient vectors, they are separated into five categories. The first category collects vectors without a dominant direction. The other categories include vectors presenting horizontal, vertical, diagonal and anti-diagonal directions. Five corresponding regularization matrices are listed in Eq. 5.

$$M_0 = \begin{pmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{pmatrix}, M_1 = \begin{pmatrix} 0 & 0 & 0 \\ -1 & 2 & -1 \\ 0 & 0 & 0 \end{pmatrix}, M_2 = \begin{pmatrix} 0 & -1 & 0 \\ 0 & 2 & 0 \\ 0 & -1 & 0 \end{pmatrix},$$
$$M_3 = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & -1 \end{pmatrix}, M_4 = \begin{pmatrix} 0 & 0 & -1 \\ 0 & 2 & 0 \\ -1 & 0 & 0 \end{pmatrix}. \quad (5)$$

They are designed to the five categories of vectors and recover edges and smooth noise in local regions. The rows of the gradient-based regularization matrix correspond to $M_j (j = 0, \ldots, 4)$ in Eq. 5 and which one is chosen depends on the local gradient vector. Horizontal and vertical vectors choose M_1 and M_2 and diagonal and anti-diagonal vectors require M_3

and M_4. Vectors without obvious orientation make use of M_0. Because f is a raster scan of image matrix, the regularization matrix C is organized in such a way that ith row of C executes an operation on the ith element of f. When $M_j (j = 0, \ldots, 4)$ is chosen to configure a certain row of C, elements of M_j are reorganized to substitute into C so the operation of M_j can be done by a multiplication of a row of C with f, instead of by a convolution.

To restore edges accurately, not only are particular regularization matrices used, but also the regularization parameter is chosen appropriately to control the degree of regularization on certain edges. Unlike the uniform regularization parameter in Tikhonov regularization and HQR, here we propose to represent the parameter as a vector, the element in which is used as a local parameter to determine the extent of regularization on a local area. ith element weights the regularization on f_i, ith element in f. Its configuration is inversely proportional to the magnitude of the gradient vector field and estimated by the expression:

$$\alpha_i = \alpha_{\max} - k\left(|\nabla f_i| - |\nabla f|_{\min}\right) \frac{\alpha_{\max} - \alpha_{\min}}{|\nabla f|_{\max} - |\nabla f|_{\min}}. \tag{6}$$

where $|\nabla f_i|$ stands for the magnitude of gradient vector at f_i and $|\nabla f|_{\min}$ and $|\nabla f|_{\max}$ are the minimal and maximal values among the gradient magnitures, respectively. k controls the rate of exponential decrease, and $\alpha_{\min}$ and $\alpha_{\max}$ are the minimal and maximal values of the parameter, which are evaluated by experiments. The design of α vector is different from the existing methods such as the discrepancy principle, generalized cross-validation and the L-curve (Kilmer and O'Leary 2001). Local gradient information is utilized to determine where and to what degree of regularization should be imposed, which is advantageous at both restoring local edges and suppressing noise according to local information.

Finally, the reconstructed solution using the nonlinear gradient-based regularization can be expressed as

$$\begin{aligned} f_{\text{reg}} &= \arg\min_f \left(\|y - Hf\|_2^2 + \|diag(\alpha)Cf\|_2^2\right), \\ &\text{subject to } f \geq 0. \end{aligned} \tag{7}$$

$diag(\alpha)$ is a diagonal matrix, whose ith diagonal element is computed by Eq. 6. C is the gradient-based regularization matrix.

4 Results

Two sets of experiments are employed to evaluate the performance of gradient-based regularization in superresolution reconstruction of images. Simulation experiment tests a reconstruction using gradient-based regularization and demonstrates its performance on edge preservation and noise suppression. The second experiment is taken on five images to show the significant improvement of gradient-based regularization on edge preservation and the reconstruction efficiency. It provides comparisons with Tikhonov regularization and HQR methods with regard to mean square error (MSE) and time consumption.

Simulation experiment is performed on a representative image, which is useful to analyze performance of reconstructions on edges. It is synthetic and of size 256×256, shown in Fig. 2a. The HR image is shifted and downsampled to generate four LR (128×128) images. White Gaussian noise corrupts these images so that SNR of them stays around 28 dB. Taking the four images as input, gradient-based regularization reconstructs a 256×256 image following the way shown in Fig. 1 and illustrated in Sect. 3. To evaluate improvements of the

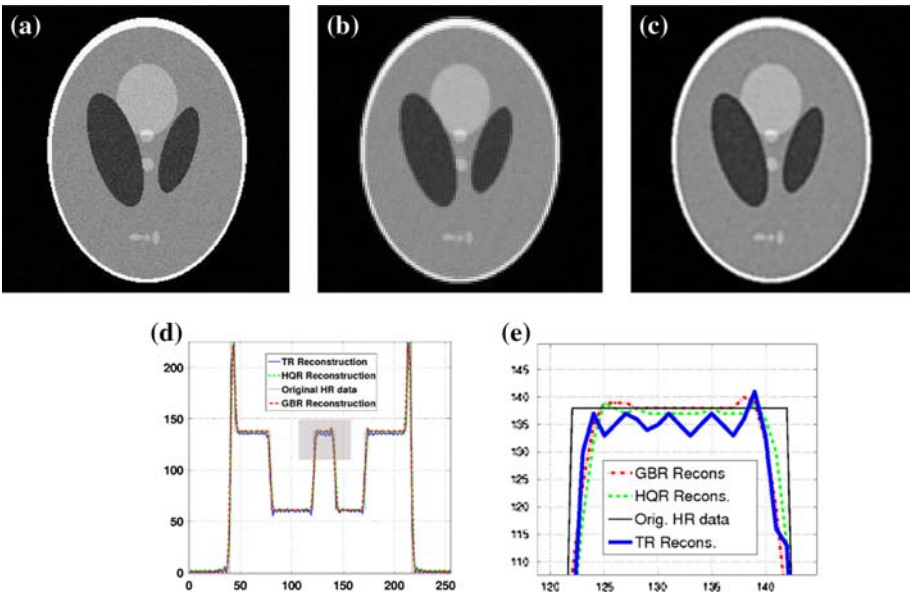

Fig. 2 Images in simulation experiment. Shown are the synthetic HR image (**a**) (the Shepp-Logan phantom commonly used in medical imaging), preliminary image (**b**) and reconstructed HR image (**c**). **d** is the 128th row of synthetic image (solid line) and reconstructed images (dashed lines) and **e** is to zoom in the gray part of (**d**)

reconstruction, the reconstructed image is analyzed in terms of MSE and time consumption and results are compared with Tikhonov regularization and HQR.

In HQR, the potential function φ takes the form

$$\varphi(\nabla f_i, s_i) = Q(\nabla f_i, s_i) + \psi(s_i). \tag{8}$$

where for every s_i the function $Q(\cdot, s_i)$ is quadratic. Both $Q(\nabla f_i, s_i)$ and $\psi(s_i)$ can be set up into either multiplicative and additive forms (Nikolova and Ng 2005). We take the additive form that can be computed faster than the multiplicative form. Hence,

$$\begin{aligned} Q(\nabla f_i, s_i) &= (\nabla f_i - s_i)^2, \\ \psi(s_i) &= |s_i|. \end{aligned} \tag{9}$$

Equation 9 has been verified to satisfy conditions of convergence and edge preservation in (Charbonnier et al. 1997; Nikolova and Ng 2001). The configuration involves two variables. They both are set up in an iteration n. We can firstly search for a reasonable value of s by $s^{n+1} = \arg\min_s(J(\nabla f^n, s))$ and then decide f by $f^{n+1} = \arg\min_f(J(\nabla f, s^{n+1}))$. The second step is possible thanks to the property of quadratic form. By alternately minimizing over f and s, the regularization converges to an edge-preserved solution for the ill-posed problem.

Gradient-based regularization follows the procedure mentioned in Sect. 3. The four LR images are organized to obtain the preliminary HR image. Sobel operators acts on the image to calculate the gradient vector field. The preliminary HR image is shown in Fig. 2b. Figure 2c displays the reconstructed HR image.

Table 1 summarizes results of MSE and time consumption of the three regularization techniques. MSE is calculated between reconstructed images and the original HR image.

Table 1 Simulation experiment results

	TR	HQR	GBR
MSE	16.03	14.03	13.86
Time (s)	23.48	984.10	28.18

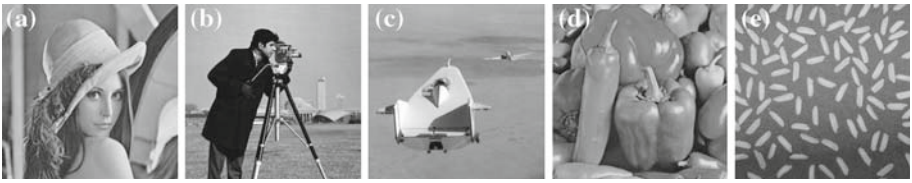

Fig. 3 Images of reconstruction experiments. Shown are five images included. **a** is referred to as lena, **b** is cameraman, **c** is liftingbody, **d** is peppers and **e** is rice

Table 2 Reconstructions comparison on five images

Images	MSE_{TR}	MSE_{HQR}	MSE_{GBR}
Lena	158.87	73.09	61.81
Cameraman	34.96	35.35	33.97
Liftingbody	37.16	41.57	27.53
Peppers	7.89	10.22	7.80
Rice	162.75	35.35	36.68

Time is counted by MATLAB. In the experiment, Gradient-based regularization presents a HR image with the least MSE. Its reconstruction time is not the shortest, but fairly equal to the time that Tikhonov regularization takes, which runs the regularized reconstruction in the shortest time. MSE evaluates the global difference between reconstructed images. The conclusion can also be verified by comparing rows of reconstructed image empirically. The 128th row of reconstructed HR image are involved for the comparison shown in Fig. 2d. Dashed line (gradient-based regularization) and dash-dot line (HQR) are almost the same and closely equal to solid line (original HR data). Taking both efficiency and performance into consideration, the reconstruction using gradient-based regularization is better at saving time and preserving edge.

Experiment is also performed on five different images shown in Fig. 3. The images are referred to as lena, cameraman, liftingbody, peppers and rice, respectively. They contain different amounts and shapes of edges. Both lena and rice have more edges, while liftingbody has less. The experiment aims to analyze performance of the reconstruction concerning edges on a variety of images. Both MSE and time consumption are measured to manifest the accuracy and efficiency. The original HR images all are the size of 256×256. The iterative process in the reconstruction by the three methods includes two stop criteria, the number of iterations and relative error. As experiments verify, five iterations are appropriate to obtain an acceptable solution.

Table 2 shows MSE comparison of the three reconstructions using Tikhonov, half-quadratic and gradient-based regularization. Before taking images into reconstruction, white Gaussian noise with zero mean and 0.001 variance is also added.

An examination of Table 2 in terms of MSE reveals that MSE of reconstruction using Tikhonov regularization is smaller when performed on peppers and much bigger on lena and rice. MSE using HQR is smaller on peppers and bigger on lena. Other MSE results vary around 35. MSE regarding gradient-based regularization is also smaller on peppers and bigger on

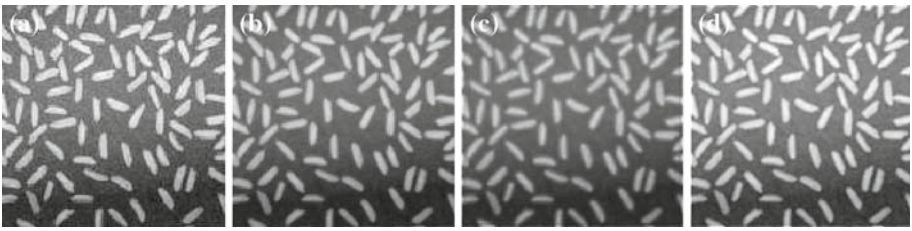

Fig. 4 Images in the second experiment. Shown are the LR image of rice (**a**), reconstructed HR images by Tikhonov regularization in (**b**), half-quadratic regularization in (**c**) and GBR method in (**d**)

lena. Compare MSE results using the three regularizations with each other. Tikhonov regularization produces MSE mostly bigger than other two methods. Gradient-based regularization presents the smaller MSE, but still bigger than that using gradient-based regularization.

So for images abound with smooth regions, such as cameraman, liftingbody and peppers, the three regularizations have similar results. But when dealing with images rich of sharp edges, such as lena and rice, gradient-based regularization shows its outstanding advantage in preserving edges accurately. Images involved in the reconstruction of rice have been included into Fig. 4.

In terms of computational time, for the series of 256×256 images, the reconstruction using Tikhonov regularization takes around 15 s with a MATLAB implementation. The gradient-based regularization needs 27 s, more than that of Tikhonov regularization by 12 s, but far less than that of HQR approximately equal to 1,000 s. Therefore, gradient-based regularization preserves edges not only accurately, but also efficiently.

5 Conclusion

In the paper, we propose a nonlinear gradient-based regularization and examine the performance on the superresolution reconstruction of images. To illustrate improvements of the regularization technique, Tikhonov regularization and HQR are studied and compared in the superresolution reconstruction of synthetic and real images. Experiments demonstrates the significant advantage of gradient-based regularization in edge preservation and time consumption. The outstanding performance can be explained by good utilization of the gradient vector field. It contains the information of local dominant orientation and intensity difference of the preliminary HR image. The other reason lies in the fact that regularization matrix and parameter are set up locally, not uniformly as in Tikhonov regularization. The assignment regularizes individual pixel according to its neighboring pixels. Therefore, gradient-based regularization can optimize the local regularization to preserve edges which is very valuable in image applications.

Acknowledgments This work was supported in part by the University Research Committee of the University of Hong Kong under Grant Number URC-10207440 and Research Grants Council of HKSAR, China (7139-06).

References

Blu, T., Bay, H., & Unser, M. (2002). A new high-resolution processing method for the deconvolution of optical coherence tomography signals. In *IEEE International Symposium on Biomedical Imaging*, Vol. 3, pp. 777–780.

Cao, N., Nehorai, A., & Jacob, M. (2007). Image reconstruction for diffuse optical tomography using sparsity regularization and expectation-maximization algorithm. *Optics Express, 15*(21), 13695–13708.

Chan, W.-S., Lam, E. Y., Ng, M. K., & Mak, G. Y. (2007). Super-resolution reconstruction in a computational compound-eye imaging system. *Multidimensional Systems and Signal Processing, 18*(2–3), 83–101.
Charbonnier, P., Blanc-Féraud, L., Aubert, G., & Barlaud, M. (1997). Deterministic edge-preserving regularization in computed imaging. *IEEE Transactions on Image Processing, 6*(2), 298–311.
Deriche, R., Kornprobst, P., Nikolova, M., & Ng, M. K. (2003). Half-quadratic regularization for MRI image restoration. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, Vol. 6, pp. 585–588.
Elad, M., & Feuer, A. (1997). Restoration of single super-resolution image from several blurred, noisy and down-sampled measure images, *IEEE Transactions on Image Processing, 6*(12), 1646–1658.
Geman, D., & Yang, C. (1995). Nonlinear image recovery with half-quadratic regularization. *IEEE Transactions on Image Processing, 4*(7), 932–946.
Kennedy, J. A., Israel, O., Frenkel, A., Bar-Shalom, R., & Azhari, H. (2006). Super-resolution in PET imaging. *IEEE Transactions on Medical Imaging, 25*(2), 137–147.
Kilmer, M. E., & O'Leary, D. P. (2001). Choosing regularization parameters in iterative methods for Ill-posed problems. *SIAM Journal on Matrix Analysis and Applications, 22*(4), 1204–1221.
Lam, E. Y. (2003). Noise in superresolution reconstruction. *Optics Letters, 28*(12), 2234–2236.
Lee, S. H., Cho, N. I., & Park, J.-I. (2003). Directional regularisation for constrained iterative image restoration. *Electronics Letters, 39*(23), 1642–1643.
Mico, V., Zalevsky, Z., & García, J. (2006). Superresolution optical system by common-path interferometry. *Optics Express, 14*(12), 5168–5177.
Ng, M. K., Shen, H., Lam, E., & Zhang, L. (2007). A total variation regularization based super-resolution reconstruction algorithm for digital video. *EURASIP Journal on Advances in Signal Processing, 2007*, 1–16, Article ID 74585.
Nikolova, M., & Ng, M. K. (2001). Fast image reconstruction algorithms combining half-quadratic regularization and preconditioning. In *Proceedings of International Conference on Image Processing*, Vol. 1, pp. 277–280.
Nikolova, M., & Ng, M. K. (2005). Fast image reconstruction algorithms combining half-quadratic regularization and preconditioning. *SIAM Journal on Scientific Computing, 27*, 937–966.
Park, S. C., Park, M. K., & Kang, M. G. (2003). Super-resolution image reconstruction: A rechnical overview. *IEEE Signal Processing Magazine, 20*(3), 21–36.
Peled, S., & Yeshurun, Y. (2001). Superresolution in MRI: Application to human white matter fiber tract visualization by diffusion tensor imaging. *Magnetic Resonance Imaging, 45*(1), 29–35.
Shin, J.-H., Sun, Y., Joung, W.-C., Paik, J.-K., & Abidi, M. A. (2001). Adaptive regularized noise smoothing of dense range image using directional Laplacian operators. In *Proceedings of SPIE Three-Dimensional Image Capture and Applications IV*, Vol. 4298, pp. 119–126.
Vanzella, W., Pellegrino, F. A., & Torre, V. (2004). Self-adaptive regularization. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 26*(6), 804–809.
Wang, Y., Hu, J., & Schroder, H. (2005). A gradient based weighted averaging method for estimation of fingerprint orientation fields. In *Proceedings of Digital Image Computing: Techniques and Applications*, p. 29.
Watzenig, D., Brandstätter, B., & Holler, G., (2004). Adaptive regularization parameter adjustment for reconstruction problems. *IEEE Transactions on Magnetics, 40*(2), 1116–1119.

Author Biographies

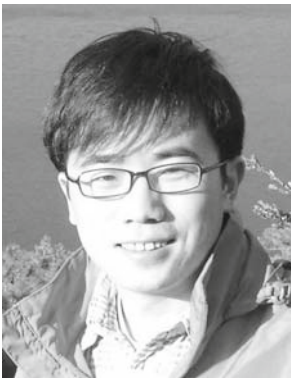

Xin Zhang received his Bachelor of Engineering degree from the Capital University of Medical Sciences and Master of Engineering degree from the Tsinghua University in 2002 and 2006, respectively. Now he is pursuing his Ph.D. degree in the University of Hong Kong. His research interests are in the areas of image reconstruction and solving inverse imaging problem.

Edmund Y. Lam received the B.S. (with distinction), M.S., and Ph.D. degrees in Electrical Engineering from Stanford University. He is now an Associate Professor in Electrical and Electronic Engineering at the University of Hong Kong, and the Founding Director of its Imaging Systems Laboratory. His research interests include computational optics and imaging, particularly their applications in semiconductor manufacturing and biomedical systems. He is also an Associate Editor of IEEE Transactions on Biomedical Circuits and Systems and has twice been a Guest Editor of the Multidimensional Systems and Signal Processing journal. He is a senior member of IEEE, and a member of OSA and SPIE.