

HESSIAN-BASED REGULARIZATION FOR 3-D MICROSCOPY IMAGE RESTORATION

Stamatios Lefkimmatis, Aurélien Bourquard, and Michael Unser

Biomedical Imaging Group, EPFL, CH-1015 Lausanne, Switzerland

Email: [stamatis.lefkimmatis,aurelien.bourquard,michael.unser]@epfl.ch

ABSTRACT

We investigate a non quadratic regularizer that is based on the Hessian operator for dealing with the restoration of 3-D images in a variational framework. We show that the regularizer under study is a valid extension of the total-variation (TV) functional, in the sense that it retains its favorable properties while following a similar underlying principle. We argue that the new functional is well suited for the restoration of 3-D biological images since it does not suffer from the well-known staircase effect of TV. Furthermore, we present an efficient 3-D algorithm for the minimization of the corresponding objective function. Finally, we validate the overall proposed regularization framework through image deblurring experiments on simulated and real biological data.

Index Terms— 3-D image restoration, Hessian matrix, Frobenius norm, mixed-norm regularization.

1. INTRODUCTION

In widefield microscopy, the imaging of biological specimens is often carried out by recording focal series of 2-D images. These images are then stacked together to generate a 3-D volume. The image series acquired using this method contain in-focus features of the specimen from the focal-plane and out-of-focus features from all adjacent planes [1]. This blurring effect is worsened by the presence of random noise, which is intrinsic to the acquisition process, and results in degraded 3-D images. This severely reduces the ability to clearly distinguish fine specimen structures. To increase the resolution and enhance the quality of images, 3-D restoration can serve as a pre-processing technique that aims to cancel-out the degradations due to the optics of the acquisition system.

Image deblurring amounts to estimating an image f from the intensity measurements y . Since a widefield microscope can be modeled in intensity as a linear space-invariant system [1], the image-observation model can be formulated as

$$y = \mathcal{A}f + w, \quad (1)$$

where \mathcal{A} is a linear blurring operator, specific to the optics of the microscope, and w is the unknown noise. The recovery of f from y is an *ill-posed problem* [2], due to the presence of noise and the operator \mathcal{A} which is usually ill-conditioned or non-invertible. To obtain a reasonable estimate of f , one must thus take into account the image-formation and acquisition processes as well as any available prior information about the properties of the image to be restored.

A common estimation strategy is to form an objective function which quantifies the quality of a given estimate and has the form

$$\mathcal{J}(f) = \mathcal{J}_{\text{data}}(f) + \tau \mathcal{R}(f). \quad (2)$$

This work was supported (in part) by the Hasler Foundation and the Indo-Swiss Joint Research Program.

The first term is known as *data fidelity* and measures the consistency between the estimate and the measurements, while the second one is the *regularization* term whose role is to impose a desirable bias on the derived solution. The *regularization parameter* $\tau \geq 0$ balances the contribution of the two terms. Image deblurring can then be cast as the minimization of (2).

As has been shown in prior studies, the choice of the regularization term can play a substantial role in the quality of the restoration. Several regularization approaches have been proposed so far for image deblurring, with a state-of-the-art method being based on minimizing the total-variation (TV) semi-norm [3]. Its success and wide use for the past two decades can be mainly attributed to its ability to produce results with well-preserved and sharp edges. Moreover, its compatibility with efficient optimization methods makes it very attractive. This is a crucial factor because the restoration of 3-D images involves processing huge amounts of data. However, TV regularization cannot be considered as a universal choice irrespectively of the class of images under consideration. In particular, it has been recently shown that Hessian-based regularizers are better adapted for the restoration of biomedical images in the 2-D case [4].

In this paper, we extend our prior work [4] to deal with the problem of 3-D image restoration. Specifically, we propose in Section 2 a higher-order non quadratic functional that is a valid extension of the TV semi-norm. The proposed regularizer is based on the Hessian Frobenius-norm and retains most of the favorable properties of TV, while it can better model the properties of biological images without introducing the undesired staircasing effect of TV. In Section 3, we devise a fast algorithm whose structure is optimized for the 3-D problem. Finally, we perform in Section 4 experiments where our approach is compared with TV on both synthetic and real measurements. We then discuss our results and the relevance of our approach for biomedical 3-D data.

2. HIGHER-ORDER REGULARIZATION

Let f be a continuously differentiable 3-D image. The TV semi-norm of f is then defined as

$$\text{TV}(f) = \int_{\Omega} \|\nabla f(\mathbf{x})\|_2 \, d\mathbf{x}, \quad (3)$$

where ∇f is the gradient of the image, $\|\nabla f\|_2$ is its Euclidean norm, and $\Omega \subset \mathbb{R}^3$. It is well-known that the first-order derivative of f in the direction specified by the unit-norm vector \mathbf{u} is given by $D_{\mathbf{u}}f = \langle \nabla f, \mathbf{u} \rangle$. With the aid of directional derivatives and by representing \mathbf{u} in spherical coordinates as

$$\mathbf{u} = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta), \quad \theta \in [0, \pi], \phi \in [0, 2\pi],$$

we can equivalently express TV as

$$\text{TV}(f) = \frac{1}{q} \int_{\Omega} \|D_{(\theta, \phi)} f(\mathbf{x})\|_{L_2(S)} \, d\mathbf{x}, \quad (4)$$

where $q = \|\sin \theta \cos \phi\|_{L_2(\mathcal{S})}$, $\mathcal{S} = [0, \pi] \times [0, 2\pi)$, and $D_{(\theta, \phi)}$ is the first directional derivative expressed in terms of the polar angle θ and the azimuthal angle ϕ .

According to (4), TV can be interpreted as a mixed L_1 - L_2 norm where the L_1 -norm acts on the image domain while the L_2 -norm acts on the domain specified by the angles (θ, ϕ) of the directional derivative. As shown below, this new interpretation of the TV functional permits us to extend its definition to higher-order differential operators. Indeed, a straightforward way to extend TV so as to include higher-order differential operators is by replacing $D_{(\theta, \phi)} f(\mathbf{x})$ in (4) by a higher-order directional derivative.

Following this approach, our work focuses on a second-order extension of TV. Accordingly, we define our new functional as

$$\mathcal{R}(f) = \frac{1}{q^2} \int_{\Omega} \|D_{\mathbf{u}, \mathbf{v}}^2 f(\mathbf{x})\|_{L_2(\mathcal{S}^2)} d\mathbf{x}, \quad (5)$$

where $D_{\mathbf{u}, \mathbf{v}}^2$ is the second-order directional derivative on the domain specified by the unit-norm vectors \mathbf{u} and \mathbf{v} . The second-order directional derivative is defined as $D_{\mathbf{u}, \mathbf{v}}^2 f(\mathbf{x}) = D_{\mathbf{u}}(D_{\mathbf{v}} f)(\mathbf{x}) = \mathbf{u}^T \mathcal{H}_f(\mathbf{x}) \mathbf{v}$, where \mathcal{H}_f is the Hessian matrix of f expressed as

$$\mathcal{H}_f = \begin{pmatrix} f_{xx} & f_{xy} & f_{xz} \\ f_{yx} & f_{yy} & f_{yz} \\ f_{zx} & f_{zy} & f_{zz} \end{pmatrix}, \quad (6)$$

with $f_{ij}(\mathbf{x}) = \frac{\partial^2}{\partial i \partial j} f(\mathbf{x})$.

Interestingly, as proven in Proposition 1, the functional in (5) is equivalent to

$$\mathcal{R}(f) = \int_{\Omega} \|\mathcal{H}_f(\mathbf{x})\|_F d\mathbf{x}, \quad (7)$$

where the integrand corresponds to the Frobenius-norm of the Hessian of f at coordinates \mathbf{x} .

Proposition 1. *The L_2 -norm of the second-order directional derivative of f at coordinates \mathbf{x} is proportional to the Frobenius-norm of the Hessian matrix $\|\mathcal{H}_f(\mathbf{x})\|_F$.*

Proof. The second-order directional derivative of f can be written as a function of the Hessian eigenvalues. Specifically, since the Hessian matrix is symmetric, we use the spectral-decomposition theorem and we express the second-order directional derivative as

$$\begin{aligned} D_{\mathbf{u}, \mathbf{v}}^2 f(\mathbf{x}) &= \mathbf{u}^T \mathbf{Q} \Lambda_f(\mathbf{x}) \mathbf{Q}^T \mathbf{v} \\ &= \left(\mathbf{Q}^T \mathbf{u} \right)^T \Lambda_f(\mathbf{x}) \left(\mathbf{Q}^T \mathbf{v} \right), \end{aligned} \quad (8)$$

where Λ_f is a 3×3 diagonal matrix with the eigenvalues of the Hessian matrix \mathcal{H}_f at coordinates \mathbf{x} and \mathbf{Q} is a 3×3 orthonormal matrix with the corresponding eigenvectors in its columns. Next we show that we can always obtain such a decomposition with \mathbf{Q} being a rotation matrix.

In order \mathbf{Q} to be a rotation matrix, it should satisfy the following two conditions:

$$(a) \mathbf{Q}^T \mathbf{Q} = \mathbf{I} \quad \text{and} \quad (b) \det(\mathbf{Q}) = 1.$$

It is clear that condition (a) is satisfied since \mathbf{Q} is by definition orthonormal. In addition, it is easy to show from (a) that $\det(\mathbf{Q}) = \pm 1$. In the case where $\det(\mathbf{Q}) = -1$, we can always choose $\mathbf{Q}' = -\mathbf{Q}$ with $\det(\mathbf{Q}') = 1$ and decompose \mathcal{H}_f as

$$\begin{aligned} \mathcal{H}_f &= \mathbf{Q}' \Lambda_f(\mathbf{x}) \mathbf{Q}'^T = (-\mathbf{Q}) \Lambda_f(\mathbf{x}) (-\mathbf{Q})^T \\ &= \mathbf{Q} \Lambda_f(\mathbf{x}) \mathbf{Q}^T. \end{aligned} \quad (9)$$

We therefore write (8) as

$$D_{\mathbf{u}, \mathbf{v}}^2 f(\mathbf{x}) = \mathbf{u}'^T \Lambda_f(\mathbf{x}) \mathbf{v}' = \sum_{k=1}^3 \lambda_k \mathbf{u}'_k \mathbf{v}'_k, \quad (10)$$

where \mathbf{u}' and \mathbf{v}' correspond to rotated versions of \mathbf{u} and \mathbf{v} , respectively, and λ_k to the diagonal elements of the matrix $\Lambda_f(\mathbf{x})$. To compute the L_2 -norm of the second-order directional derivative we express the unit-vectors \mathbf{u}' and \mathbf{v}' in spherical coordinates as

$$\begin{aligned} \mathbf{u}' &= (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta), \quad \theta \in [0, \pi], \phi \in [0, 2\pi), \\ \mathbf{v}' &= (\sin \alpha \cos \beta, \sin \alpha \sin \beta, \cos \alpha), \quad \alpha \in [0, \pi], \beta \in [0, 2\pi). \end{aligned}$$

Based on (10), we then obtain

$$\begin{aligned} \|D_{\mathbf{u}, \mathbf{v}}^2 f(\mathbf{x})\|_{L_2(\mathcal{S}^2)} &= \left(\int_{\mathcal{S}^2} |D_{\mathbf{u}, \mathbf{v}}^2 f(\mathbf{x})|^2 \sin \theta \sin \alpha d\mathbf{V} \right)^{1/2} \\ &= \sqrt{\frac{16\pi^2}{9} (\lambda_1^2 + \lambda_2^2 + \lambda_3^2)} \\ &= \frac{4\pi}{3} \|\mathcal{H}_f(\mathbf{x})\|_F, \end{aligned} \quad (11)$$

where $d\mathbf{V} = d\theta d\phi d\alpha d\beta$. \square

From Definition (7), we can see that the proposed functional corresponds to a convex regularizer since it arises from the integration of a linear operator [5]. Moreover, it is easy to verify that our regularizer also retains the important properties of TV, namely, homogeneity, rotation, and translation invariance.

Since the Frobenius norm of a matrix is equal to the Euclidian norm of its vectorized version, we can also write (7) as

$$\mathcal{R}(f) = \int_{\Omega} \|\mathcal{U}f(\mathbf{x})\|_2 d\mathbf{x}, \quad (12)$$

where \mathcal{U} is a differential operator defined as $\mathcal{U} = \{\partial_{ii}, \sqrt{2}\partial_{ij}\}_{i \neq j}$. This reformulation of the Hessian Frobenius-norm regularizer is advantageous for the description of the minimization algorithm we present in Section 3.

3. MINIMIZATION OF THE OBJECTIVE FUNCTION

3.1. Majorization-Minimization Approach

In this section, we depart from the continuous domain and we consider the discrete formulation of the image-restoration problem. Assuming that additive Gaussian noise is degrading the measurements, the appropriate data term is quadratic. Thus, the overall objective function reads as

$$\mathcal{J}(\mathbf{f}) = \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{f}\|_2^2 + \tau \mathcal{R}(\mathbf{f}), \quad (13)$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the convolution matrix describing the blurring operation and $\mathbf{y}, \mathbf{f} \in \mathbb{R}^N$ are the N -dimensional rasterized observed and unknown images, respectively, with $N = n_1 \times n_2 \times n_3$.

Defining the ℓ_1 norm of a vector field $\mathbf{u} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N) \in \mathbb{R}^{N \times k}$ as $\|\mathbf{u}\|_1 = \sum_{1 \leq i \leq N} |\mathbf{u}_i|$, where $|\mathbf{u}_i| = \left(\sum_{1 \leq j \leq k} \mathbf{u}_{i,j}^2 \right)^{1/2}$, we re-write the proposed regularizer in its discrete form as

$$\mathcal{R}(\mathbf{f}) = \|\mathbf{U}\mathbf{f}\|_1 = \sum_{i=1}^N |(\mathbf{U}\mathbf{f})_i|, \quad (14)$$

where $(\cdot)_i$ denotes the i th element of the argument.

We minimize (13) following a *majorization-minimization* (MM) approach [4, 6, 7]. To develop a MM-based algorithm, we first derive an appropriate quadratic majorizer $\mathcal{Q}_R(\mathbf{f}; \mathbf{f}')$ of our penalty function $\mathcal{R}(\mathbf{f})$. To do so we employ the inequality

$$\sqrt{g(x)} \leq \frac{\sqrt{g(y)}}{2} + \frac{g(x)}{2\sqrt{g(y)}}, \quad (15)$$

that holds true for any general function $g(\cdot) : \mathbb{R} \mapsto \mathbb{R}$, $\forall (x, y) : g(x) \geq 0, g(y) > 0$, with equality in (15) if and only if $g(x) = g(y)$. Using the property that the majorization relation is closed under the formation of sums and non-negative products [7], we can then show that the function

$$\mathcal{Q}_R(\mathbf{f}; \mathbf{f}^{(t)}) = \frac{1}{2} \|\mathbf{U}\mathbf{f}^{(t)}\|_1 + \frac{1}{2} \sum_{i=1}^N \frac{|\langle \mathbf{U}\mathbf{f} \rangle_i|^2}{|\langle \mathbf{U}\mathbf{f}^{(t)} \rangle_i|} \quad (16)$$

is a valid majorizer of $\mathcal{R}(\mathbf{f})$ at the fixed point $\mathbf{f}^{(t)}$.

Since the data-fidelity term $\mathcal{J}_{\text{data}}(\mathbf{f})$ in (13) is quadratic, the function $\mathcal{Q}(\mathbf{f}; \mathbf{f}') = \mathcal{J}_{\text{data}}(\mathbf{f}) + \mathcal{Q}_R(\mathbf{f}; \mathbf{f}')$ is itself a quadratic majorizer of the complete objective function $\mathcal{J}(\mathbf{f})$. This implies that the minimization of the resulting majorizer amounts to solving the system of linear equations

$$\underbrace{(\mathbf{A}^\top \mathbf{A} + \tau \mathbf{U}^\top \mathbf{W}^{(t)} \mathbf{U})}_{\mathbf{S}^{(t)}} \mathbf{f}^{(t+1)} = \mathbf{A}^\top \mathbf{y}, \quad (17)$$

where $\mathbf{W}^{(t)}$ is an $N \times N$ block-diagonal matrix with diagonal components

$$(\mathbf{W}^{(t)})_{ii} = \frac{1}{2|\langle \mathbf{U}\mathbf{f}^{(t)} \rangle_i|}, i \in [1, N]. \quad (18)$$

To solve (17), we then employ an efficient preconditioned conjugate-gradient (PCG) method [8] with a diagonal preconditioner. Specifically, we use a left-and-right preconditioning scheme that exploits the diagonal part of $\mathbf{S}^{(t)}$, which is most efficient in our 3-D setting since no FFT is involved.

We describe now how these diagonal-matrix components can be determined. Let us consider the generic matrix structure $\mathbf{M} = \mathbf{B}^\top \mathbf{D} \mathbf{B}$, where \mathbf{B} and \mathbf{D} are $N \times N$ circulant and diagonal matrices, respectively. Given a rasterized image vector \mathbf{c} , the product $\mathbf{c}' = \mathbf{M}\mathbf{c}$ in vector notation corresponds to the discrete spatial-domain operation

$$c' = \sum_{\mathbf{n}_i} b^\top[\mathbf{n}_2] b[\mathbf{n}_1] d[\cdot - \mathbf{n}_2] c[\cdot - (\mathbf{n}_1 + \mathbf{n}_2)], \quad (19)$$

where c, c' are the corresponding image sequences and b, d are the filter and the pointwise-multiplication map associated to the matrices \mathbf{B}, \mathbf{D} , respectively. Based on (19), we can show that the expression of the diagonal part of \mathbf{M} corresponds to the spatial-domain pointwise-multiplication map

$$\mathcal{D}[\mathbf{k}] = [d \star (b \bullet b)^\top][\mathbf{k}], \quad (20)$$

where \star and \bullet denote discrete convolution and pointwise multiplication, respectively. From its definition, $\mathbf{S}^{(t)}$ decomposes as a sum of matrices of the same structure as \mathbf{M} . From (20), and by linearity, the diagonal part of the whole system matrix can thus be determined. Note that the $\mathbf{A}^\top \mathbf{A}$ term appearing in (17) can be handled in (20) as a degenerate case where d reduces to the identity.

Table 1. PSNR comparisons of TV and Hessian Frobenius-norm regularization on two test stacks of images.

Data	Stack A			Stack B		
Methods	Blurred	TV	Frob.	Blurred	TV	Frob.
BSNR 15 dB	23.93	26.19	26.35	27.92	30.51	30.58
BSNR 20 dB	24.34	26.65	26.86	28.27	30.89	30.97

3.2. Relation between MM and Lagged-Diffusivity

One of the standard approaches for minimizing (13) under TV regularization is the method of *lagged diffusivity* that was first proposed in [9]. This method is based on the Euler-Lagrange equation and leads to an elliptical partial differential equation (PDE). To solve this PDE a fixed-point iterative approach was presented in [9]. The global convergence of its discrete approximations was established in [10].

Following a similar approach for our continuous regularizer (12), the corresponding Euler-Lagrange equation leads to the fourth-order PDE

$$g(f) = \mathcal{A}^*(\mathcal{A}f - y) + \tau \mathcal{U} \cdot \left(\frac{\mathcal{U}f}{\|\mathcal{U}f\|_2} \right). \quad (21)$$

The global minimizer of the objective function is then found as a sequence of solutions to the system of equations

$$\mathcal{A}^* \mathcal{A}f^{(t+1)} + \tau \mathcal{L}(f^{(t)}) f^{(t+1)} = \mathcal{A}^* y, \quad (22)$$

where \mathcal{A}^* is the adjoint of the blurring operator \mathcal{A} and $\mathcal{L}(f)$ is an operator whose action on z is given by

$$\mathcal{L}(f)z = \mathcal{U} \cdot \left(\frac{\mathcal{U}z}{\|\mathcal{U}z\|_2} \right). \quad (23)$$

This minimization strategy earns the name of *lagged diffusivity* because, in order to obtain the next iterate $f^{(t+1)}$ that closer approximates the solution to the minimization problem, one has to solve a PDE whose diffusivities are fixed and depend on the previous iterate $f^{(t)}$. Upon discretization of the solution, this iterative approach coincides with the finite MM method we have proposed in Section 3.1.

4. EXPERIMENTS

To validate the effectiveness of the proposed Hessian-based regularization framework, we provide experimental results for the task of 3-D image deblurring. We compare our results with TV on two stacks of images. In our experiments, we use a Gaussian *point-spread function* (psf) of support $20 \times 20 \times 10$. The standard deviation for the Gaussian PSF changes linearly in the z-axis from $\sigma_a = 2$ to $\sigma_b = 10$. In Table 1, we provide numerical results in terms of PSNR for two BSNRs ($\text{BSNR} = \text{var}[\mathbf{A}\mathbf{f}] / \sigma^2$) corresponding to different levels of Gaussian noise. For the sake of fairness, the results reported for each regularizer are obtained using the individualized regularization parameter τ that gives the best PSNR performance. For the discretization of the differential operators, we use forward finite differences. Finally, regarding the minimization of the objective functions, we run 10 successive quadratic-bound minimizations and for each one we use the PCG algorithm with a stopping criterion set to either reaching a relative normed difference of 10^{-5} between two successive estimates, or a maximum of 20 iterations.

Apart from the quantitative results of Table 1, the efficacy of our approach can be also appreciated visually from the representative deblurring examples shown in Figs. 1 and 2. In these examples,




Fig. 1. Restoration results on the stack A of size $256 \times 256 \times 52$. Maximum-intensity projection (along the z-axis) of the (a) blurred stack (PSNR=24.34 dB), (b) TV reconstruction (PSNR=26.65 dB), (c) Hessian Frobenius-norm reconstruction (PSNR=26.86 dB).




Fig. 2. Restoration results on the stack B of size $330 \times 330 \times 197$. Maximum-intensity projection (along the z-axis) of the (a) blurred stack (PSNR=27.92 dB), (b) TV reconstruction (PSNR=30.51 dB), (c) Hessian Frobenius-norm reconstruction (PSNR=30.58 dB).

despite the seemingly small differences in PSNR improvement between the two regularizers, we ascertain that, while TV regularization introduces heavy staircase artifacts mixing structural details of the images, our regularizer matches better the intensity variations and results in a better reconstruction. To evaluate the practical relevance of our approach, we also present in Fig. 3 deblurring results on a real 3-D fluorescence two-channel widefield-microscopy stack using a standard diffraction-limited PSF model [1]. As visual reference, we use the confocal counterpart of the acquired data. In this last example, we once more verify that—in contrast to TV—the Frobenius-norm regularizer leads to solutions that do not suffer from mixing or blocking artifacts. It can thus be considered as a better choice, especially when one has to deal with images that consist mostly of ridges and filament-like structures.

5. CONCLUSIONS

We have proposed a second-order extension of TV for the reconstruction of 3-D images. We have demonstrated that the resulting Hessian Frobenius-norm regularizer can perform better than TV for images that consist mostly of ridges and smooth transitions of intensities, both from a qualitative and quantitative point of view. In particular, our regularizer can match continuous intensity variations better. Therefore, it can circumvent the staircase effects as well fine-scale-structure deformations that occur with TV.

Acknowledgments: The authors would like to thank Daniel Sage for helpful discussions, Philippe Thévenaz for useful comments and careful proofreading, Cédric Vonesch for kindly providing the widefield and confocal stacks of Fig. 3, and George von Dassow for kindly providing the stacks used in Figs. 1–2.




Fig. 3. Restoration results on a real two-channel fluorescent-cell stack of size $352 \times 512 \times 96$. Maximum-intensity projection (along the z-axis) of the (a) widefield stack, (b) reference confocal stack, (c) TV reconstruction, (d) Hessian Frobenius-norm reconstruction. The details of this figure are better seen in the electronic version of this paper by zooming on the screen.

6. REFERENCES

- [1] C. Vonesch, F. Aguet, J.-L. Vonesch, and M. Unser, “The colored revolution of bioimaging,” *IEEE Sig. Process. Magazine*, vol. 23, no. 3, pp. 20–31, 2006.
- [2] P. C. Hansen, J. G. Nagy, and D. P. O’Leary, *Deblurring Images: Matrices, Spectra, and Filtering*, SIAM, 2006.
- [3] L. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D*, vol. 60, pp. 259–268, 1992.
- [4] S. Lefkimmiatis, A. Bourquard, and M. Unser, “Hessian-based norm regularization for image restoration with biomedical applications,” *IEEE Trans. Image Processing*, to appear.
- [5] S. Boyd and L. Vandenberghe, *Convex Optimization*, Kluwer Academic Publishers, 2004.
- [6] D. Hunter and K. Lange, “A tutorial on MM algorithms,” *The American Statistician*, vol. 58, pp. 30–37, 2004.
- [7] M.A.T. Figueiredo, J.M. Bioucas-Dias, and R.D. Nowak, “Majorization–minimization algorithms for wavelet-based image restoration,” *IEEE Trans. Image Processing*, vol. 16, pp. 2980–2991, 2007.
- [8] J.R. Shewchuk, “An introduction to the conjugate gradient method without the agonizing pain,” 1994, [Online] Available: <http://www.cs.cmu.edu/~jrs/jrspapers.html>.
- [9] C.R. Vogel and M.E. Oman, “Iterative methods for total variation denoising,” *SIAM J. Sci. Comput.*, vol. 17, pp. 227–238, 1996.
- [10] D.C. Dobson and C.R. Vogel, “Convergence of an iterative method for total variation denoising,” *SIAM J. Numer. Anal.*, vol. 34, pp. 1779–1791, 1997.