

A FAST THRESHOLDED LANDWEBER ALGORITHM FOR GENERAL WAVELET BASES: APPLICATION TO 3D DECONVOLUTION MICROSCOPY

Cédric Vonesch and Michael Unser

Biomedical Imaging Group, EPFL, Lausanne, Switzerland

ABSTRACT

Wavelet-domain ℓ_1 -regularization is a promising approach to deconvolution. The corresponding variational problem can be solved using a “thresholded Landweber” (TL) algorithm. While this iterative procedure is simple to implement, it is known to converge slowly. In this paper, we give the principle of a modified algorithm that is substantially faster. The method is applicable to arbitrary wavelet representations, thus generalizing our previous work which was restricted to the orthonormal Shannon wavelet basis.

Numerical experiments show that we can obtain up to a 10-fold speed-up with respect to the existing TL algorithm, while providing the same restoration quality. We also present an example with real data that demonstrates the feasibility of wavelet-domain regularization for 3D deconvolution microscopy.

Index Terms— Deconvolution, fast, wavelets, sparsity, ℓ_1 -minimization, 3D, microscopy.

1. INTRODUCTION

Modern experimental biology makes extensive use of fluorescent probes for selectively labeling structures of interest [1]. In this context, optical microscopy plays a key role. An interesting challenge arises when the structures to be observed are just a few times larger than the physical resolution limit of the imaging instrument. In such situations, it makes sense to employ numerical image-enhancement methods. When applied to widefield systems, this approach is commonly dubbed “deconvolution microscopy”. One of its main difficulties lies in the size of the data sets that are routinely produced in three-dimensional (3D) fluorescence microscopy. This puts a strong limitation on the computational complexity of potential deconvolution procedures. In particular, deconvolution methods of the latest generation are often considered too heavy computationally and are scarcely used in practice.

A relatively novel approach to deconvolution is (non-quadratic) wavelet-domain regularization, which was shown to yield significantly better results than older-generation methods [2]. The underlying variational problem can be

solved using a simple iterative algorithm [2, 3, 4]; however, this algorithm is known to converge slowly [2, 5, 6].

In previous work [6], we decoupled the problem into subbands and derived a modified algorithm that is substantially faster for the orthonormal Shannon wavelet basis. The goals of the present work are 1) to give the principle of a comparably fast algorithm that allows for more general (non-bandlimited) wavelet representations and 2) to demonstrate its efficiency on synthetic data as well as real 3D fluorescence images.

2. VARIATIONAL FORMULATION OF DECONVOLUTION USING WAVELET-DOMAIN REGULARIZATION

Image formation in a widefield fluorescence microscope can be modeled by a linear transformation that maps a 3D distribution of fluorescent markers into a 3D distribution of light intensity. Under normal optical conditions, this process is essentially shift-invariant. The light intensity is typically measured by a CCD detector. Using an appropriate discretization, the complete imaging process can thus be described by the following equation:

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{b},$$

where \mathbf{H} is a convolution (i.e. block-circulant) matrix and \mathbf{b} is a vector representing the distortions introduced by the optics and/or the measurement device. The problem of deconvolution is to estimate the signal of interest \mathbf{x} from the measurement \mathbf{y} . This estimation can be extremely unstable when \mathbf{H} is ill-conditioned—such as in deconvolution microscopy—and is thus usually performed with the help of regularization techniques [7].

Wavelet-domain regularization is based on the paradigm that many natural signals can be well approximated using only a few elements from a given wavelet basis. In our case, this means that $\mathbf{x} \simeq \mathbf{W}\mathbf{w}$, where \mathbf{W} is a “synthesis matrix” whose columns are the elements of the wavelet basis and \mathbf{w} is a vector of wavelet coefficients with a small number of non-zero entries. It is therefore natural to estimate \mathbf{x} in the wavelet domain, while imposing a “sparsity” constraint. This can be

This work was supported by the Hasler Foundation.




Fig. 1. A bivariate example showing that the cost functional $\mathcal{C}(\mathbf{w}) = \mathcal{C}(w_1, w_2)$ usually has a sparse minimizer. The ℓ_1 -sphere on which the minimizer lies is determined by the regularization parameter λ . The elliptic curves represent the level sets of the discrepancy term; the one corresponding to the smallest feasible value usually passes through a “corner” of the ℓ_1 -sphere.

done (see Fig. 1) by minimizing the cost functional

$$\mathcal{C}(\mathbf{w}) = \|\mathbf{y} - \mathbf{H}\mathbf{W}\mathbf{w}\|_2^2 + \lambda\|\mathbf{w}\|_1.$$

Here, the first term represents the squared Euclidian norm of the discrepancy between the measurement and the estimate; the second term is the sum of the absolute values of the wavelet coefficients (up to the multiplicative regularization parameter λ , which balances the relative influence of both terms).

3. THE FAST THRESHOLDED LANDWEBER ALGORITHM

Several research groups [2, 3, 4] have independently proposed to minimize $\mathcal{C}(\mathbf{w})$ using a two-step procedure that alternates between a Landweber iteration [7] and a soft-thresholding operation; hence the name “thresholded Landweber” (TL) algorithm. More precisely, starting from an arbitrary initialization $\tilde{\mathbf{w}}$, the TL algorithm consists in the recursive application of the update formula

$$\mathbf{w} = \mathcal{T}_{\lambda\tau/2}\{\tilde{\mathbf{w}} + \tau\mathbf{W}^T\mathbf{H}^T(\mathbf{y} - \mathbf{H}\mathbf{W}\tilde{\mathbf{w}})\}. \quad (1)$$

Here, τ is the step size of the Landweber iteration and $\mathcal{T}_\theta\{\cdot\}$ denotes a componentwise application of the soft-thresholding

function

$$\mathcal{T}_\theta(z) = \text{sgn}(z) \max(|z| - \theta, 0).$$

In this section, we revisit the derivation of the TL algorithm presented in [3], so as to obtain a faster procedure. A specificity of our approach is that we take into account the subband structure of the wavelet family. To do so, we introduce a set S that indexes the different subbands. The subscript s will be used to refer to a subband $s \in S$. For example, \mathbf{w}_s is the vector that contains the coefficients of \mathbf{w} corresponding to subband s and \mathbf{W}_s is the matrix that reconstructs the component corresponding to these coefficients. In particular, we can write that

$$\mathbf{W}\mathbf{w} = \sum_{s \in S} \mathbf{W}_s \mathbf{w}_s.$$

Suppose that we are given an estimate $\tilde{\mathbf{w}}$ of the minimizer of the cost functional \mathcal{C} . We then define the auxiliary functional

$$\mathcal{A}(\mathbf{w}) = \mathcal{C}(\mathbf{w}) - \|\mathbf{H}\mathbf{W}(\mathbf{w} - \tilde{\mathbf{w}})\|_2^2 + \sum_{s \in S} \alpha_s \|\mathbf{w}_s - \tilde{\mathbf{w}}_s\|_2^2,$$

where the constants α_s are real and strictly positive. We assume that they are chosen such that $\mathcal{A}(\mathbf{w}) \geq \mathcal{C}(\mathbf{w})$, which is always possible; furthermore, the above definition implies that $\mathcal{A}(\mathbf{w}) = \mathcal{C}(\tilde{\mathbf{w}})$ when $\mathbf{w} = \tilde{\mathbf{w}}$. These two properties guarantee that, if we find a \mathbf{w} such that $\mathcal{A}(\mathbf{w}) \leq \mathcal{A}(\tilde{\mathbf{w}})$, we also have $\mathcal{C}(\mathbf{w}) \leq \mathcal{C}(\tilde{\mathbf{w}})$. In other terms, if we decrease the auxiliary functional $\mathcal{A}(\mathbf{w})$, we also decrease the cost functional $\mathcal{C}(\mathbf{w})$.

It turns out that the coefficients of \mathbf{w} are completely decoupled in this auxiliary functional, so that we can obtain a closed-form expression of its minimizer. Indeed, defining $\tau_s = \alpha_s^{-1}$ and $\mathbf{z}_s = \tilde{\mathbf{w}}_s + \tau_s \mathbf{W}_s^T \mathbf{A}^T (\mathbf{y} - \mathbf{A}\mathbf{W}\tilde{\mathbf{w}})$, a short computation shows that

$$\mathcal{A}(\mathbf{w}) = c + \sum_{s \in S} \alpha_s \left(\|\mathbf{z}_s - \mathbf{w}_s\|_2^2 + \lambda \tau_s \|\mathbf{w}_s\|_1 \right), \quad (2)$$

where c is a constant with respect to \mathbf{w} . Thus, once the intermediate vectors \mathbf{z}_s have been computed, $\mathcal{A}(\mathbf{w})$ is essentially a positive linear combination of monovariate functionals of the form $|z - w|^2 + 2\theta|w|$. One can verify that the minimizer of such a monovariate functional is given by the soft-thresholding function defined above.

It follows that the minimizer of the auxiliary functional is given by $\mathbf{w}_s = \mathcal{T}_{\lambda\tau_s/2}\{\tilde{\mathbf{w}}_s + \tau_s \mathbf{W}_s^T \mathbf{H}^T (\mathbf{y} - \mathbf{H}\mathbf{W}\tilde{\mathbf{w}})\}$, for each subband s . When the constants α_s have the same value α for all subbands, this leads precisely to the TL algorithm described in the beginning, with a step-size $\tau = \alpha^{-1}$.

It is however not necessary to compute the exact minimizer of $\mathcal{A}(\mathbf{w})$. In particular, expression (2) shows that the auxiliary functional is still decreased if we compute the minimizer only for a single subband s , that is, if we define

$$\begin{cases} \mathbf{w}_s = \mathcal{T}_{\lambda\tau_s/2}\{\tilde{\mathbf{w}}_s + \tau_s \mathbf{W}_s^T \mathbf{H}^T (\mathbf{y} - \mathbf{H}\mathbf{W}\tilde{\mathbf{w}})\}; \\ \mathbf{w}_{s'} = \tilde{\mathbf{w}}_{s'} \text{ for } s' \in S \setminus \{s\}. \end{cases} \quad (3)$$

This choice allows us to adjust the constant α_s similarly to what we did in [6], while not being bound to the Shannon wavelet basis. The corresponding step size can be significantly larger than in the standard TL algorithm, thereby providing a more effective update for subband s . Thus, when applied successively to every subband, we call procedure (3) the “fast thresholded Landweber” (FTL) algorithm. The details of its implementation will be discussed in a forthcoming paper. The main difference with [6] is that the computation of the residual is more involved because the subbands are coupled for non-bandlimited wavelets.

4. EXPERIMENTAL RESULTS

We now present experimental results that confirm the efficiency of the method for various wavelet families. We generated a synthetic test case by convolving the standard 256×256 MRI image with a Gaussian kernel of width $\sigma_{\text{kernel}} = 2$ pixels, and by adding white Gaussian noise of standard deviation $\sigma_{\text{noise}} = 0.4266$ to the result. We then compared the SNR improvement as a function of time during the execution of the FTL and TL algorithms, with four classical wavelet bases. The measurements were averaged over 30 noise realizations. Both algorithms were used with the same settings. Since our purpose is not to discuss restoration quality itself, we settled on using the same value $\lambda = 0.2$ in all situations, which gave satisfying results¹ independently of the wavelet basis. Furthermore, a 3-level decomposition depth was used for all wavelet representations. We also used the same initialization for the TL and FTL algorithms (obtained using a quadratic Tikhonov regularization). Finally, we applied a random-shift method for both algorithms, as in the work of Figueiredo and Nowak [2].

Although the computation times depend on the software and hardware environment (we implemented the algorithms under Matlab on a 2.66 GHz Intel Xeon workstation), they are representative of the acceleration that can be achieved in practice. The graphs presented in Fig. 2 show the following consistent trend: to achieve an SNR improvement comprised between 7 and 8 dB, the FTL algorithm is faster than the TL algorithm by a factor that varies roughly between 5 (for the Haar basis) and 10 (for the 9/7 wavelet basis). The latter value is comparable to the acceleration factor that we obtained for the Shannon wavelet basis [6].

The FTL algorithm thus makes wavelet regularization tractable for larger-scale deconvolution problems. We applied it for deconvolving a 3D stack of real fluorescence images (260 slices of 384×448 pixels). The data was acquired on a standard widefield epifluorescence microscope with a $63 \times$ oil-immersion objective, and represents a *C. Elegans* embryo labeled with three fluorophores (Hoechst, Alexa488 and Alexa568). Its maximum-intensity projection (along the

¹The SNR figures are comparable to the ones obtained for a similar test case in [2], where a more extensive study of restoration quality is proposed.




Fig. 2. Comparison of the SNR improvement (in dB) as a function of computation time (in seconds) for the TL and FTL algorithms, using various wavelet bases.

axial dimension) is shown in Fig. 3 (top). Each channel was processed separately, using a computed PSF based on a diffraction-limited model [1] which takes into account the numerical aperture of the objective, the refractive index of the immersion medium and the emission wavelength. The deconvolution was performed with the Haar basis, using 3 decomposition levels for the radial (X-Y) dimensions and 2 decomposition levels for the axial (Z) dimension. Fig. 3 (bottom) shows the maximum-intensity projection of the result. It is seen that the typical widefield haze (due to out-of-focus light) is significantly reduced. Also, many details inside the cell are visible with more contrast (e.g. the microtubules and the chromosomes). In the red channel, the vesicles are brighter relatively to the background.




Fig. 3. Maximum intensity projections of a 3D widefield image stack (top) and its deconvolved version using the FTL algorithm (bottom).

5. CONCLUSION

We have presented an algorithm for wavelet-regularized deconvolution that is up to one order of magnitude faster than the standard method and works for arbitrary wavelet bases. This makes wavelet regularization a serious alternative for large-scale restoration problems; we have successfully applied it to 3D deconvolution microscopy.

In principle, the same algorithm could also be applied with redundant wavelet representations². Although our preliminary results do not reveal significant differences between wavelet bases, future work should explore this aspect in more details.

6. REFERENCES

- [1] C. Vonesch, F. Aguet, J.-L. Vonesch, and M. Unser, "The colored revolution of bioimaging," *IEEE Signal Processing Magazine*, vol. 23, no. 3, pp. 20–31, May 2006.
- [2] M. A. T. Figueiredo and R. D. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Transactions on Image Processing*, vol. 12, no. 8, pp. 906–916, August 2003.
- [3] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Communications on Pure and Applied Mathematics*, vol. 57, no. 11, pp. 1413–1457, August 2004.
- [4] J. Bect, L. Blanc-Féraud, G. Aubert, and A. Chambolle, "A ℓ^1 -unified variational framework for image restoration," *Lecture Notes in Computer Science - Proceedings of ECCV 2004 Part IV*, vol. 3024, pp. 1–13, 2004.
- [5] I. Loris, G. Nolet, I. Daubechies, and T. Dahlen, "Tomographic inversion using ℓ_1 -regularization of wavelet coefficients," *Geophysical Journal International*, vol. 170, no. 1, pp. 359–370, July 2007.
- [6] C. Vonesch and M. Unser, "A fast thresholded Landweber algorithm for wavelet-regularized multidimensional deconvolution," *IEEE Transactions on Image Processing*, in press.
- [7] M. Bertero and P. Boccacci, *Introduction to inverse problems in imaging*, Institute of Physics Publishing, 1998.

²However, the memory requirements would be too high for 3D deconvolution microscopy.