

An EM Algorithm for Wavelet-Based Image Restoration

Mário A. T. Figueiredo, *Senior Member, IEEE*, and Robert D. Nowak, *Member, IEEE*

Abstract

This paper introduces an *expectation-maximization* (EM) algorithm for image restoration (deconvolution) based on a penalized likelihood formulated in the wavelet domain. Regularization is achieved by promoting a reconstruction with low-complexity, expressed in the wavelet coefficients, taking advantage of the well known sparsity of wavelet representations. Previous works have investigated wavelet-based restoration but, except for certain special cases, the resulting criteria are solved approximately or require demanding optimization methods. The EM algorithm herein proposed combines the efficient image representation offered by the discrete wavelet transform (DWT) with the diagonalization of the convolution operator obtained in the Fourier domain. Thus, it is the first general-purpose approach to wavelet-based image restoration with computational complexity comparable to that of standard wavelet denoising schemes or of frequency domain deconvolution methods. The algorithm alternates between an E-step based on the fast Fourier transform (FFT) and a DWT-based M-step, resulting in an efficient iterative process requiring $O(N \log N)$ operations per iteration. The convergence behavior of the algorithm is investigated, and it is shown that under mild conditions the algorithm converges to a globally optimal restoration. Moreover, our new approach performs competitively with, in some cases better than, the best existing methods in benchmark tests.

M. Figueiredo is with the Institute of Telecommunications and the Department of Electrical and Computer Engineering, *Instituto Superior Técnico*, 1049-001 Lisboa, Portugal (E-mail: mtf@lx.it.pt).

R. Nowak is with the Department of Electrical and Computer Engineering, 2041 Duncan Hall, Rice University, Houston, TX, USA (e-mail: nowak@rice.edu).

M. Figueiredo was partially supported the Science and Technology Foundation (FCT, Portugal), grant POSI/33143/SRI/2000.

R. Nowak was partially supported by the National Science Foundation, grant no. MIP-9701692, the Army Research Office, grant no. DAAD19-99-1-0349, the Office of Naval Research, grant no. N00014-00-1-0390.

I. INTRODUCTION

Wavelet-based methods had an important impact on the field of image processing, especially in coding and denoising. Their success is supported on the fact that the wavelet transforms of images tend to be sparse (*i.e.*, most of the wavelet coefficients are close to zero). This implies that image approximations based on a small subset of wavelets are typically very accurate, which is a key to wavelet-based compression. The good performance of wavelet-based denoising is also intimately related to the approximation capabilities of wavelets. Thus, the conventional wisdom is that wavelet representations that provide good approximations will also perform well in estimation problems [23].

Image deconvolution is a more challenging problem than denoising. This is a classic, well-studied image processing task [1], but applying wavelets has proved to be a challenging problem. Deconvolution is most easily dealt with (at least computationally) in the Fourier domain. However, image modelling (thus denoising) is best handled in the wavelet domain; here lies the problem. Convolution operators are generally quite difficult to represent in the wavelet domain, unlike the simple diagonalization obtained in the Fourier domain. This naturally suggests the possibility of combining Fourier-based deconvolution and wavelet-based denoising, and several ad hoc proposals for this sort of combination have appeared in the literature.

In this paper we formally develop an image deconvolution algorithm based on a *maximum penalized likelihood estimator* (MPLE). The MPLE cannot be computed in closed-form, and so we propose an *expectation-maximization* (EM) algorithm to numerically compute it. The result is an EM iterative deconvolution algorithm which alternates between the Fourier and wavelet domains. We compare our results with state-of-the-art methods in benchmark problems and show that it performs competitively, sometimes better, in terms of the SNR improvement achieved.

II. PROBLEM FORMULATION

The goal of *image restoration* is to recover an *original image* \mathbf{x} from a degraded (or imperfect) observed version \mathbf{y} [1]. In this paper, \mathbf{x} and \mathbf{y} will denote vectors containing all the image

pixel values, after some (*e.g.*, lexicographic) ordering. Let N_x and N_y be the dimensionality of vectors \mathbf{x} and \mathbf{y} , respectively. The class of observations/degradations considered in this paper is described by the standard “linear observation plus Gaussian noise” model:

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}. \quad (1)$$

In (1), \mathbf{H} denotes the (linear) observation operator (*i.e.*, a $N_y \times N_x$ matrix), and \mathbf{n} is a sample of a zero-mean white Gaussian noise process with variance σ^2 ; that is, $p(\mathbf{n}) = \mathcal{N}(\mathbf{n}|0, \sigma^2\mathbf{I})$, where $\mathcal{N}(\mathbf{g}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes a multivariate Gaussian density with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$, evaluated at \mathbf{g} , and \mathbf{I} is an identity matrix. Typical observation mechanisms which are adequately approximated by (1) include: optical blur, motion blur, tomographic projections, refraction and/or multipath effects (*e.g.*, in underwater imaging), electronic noise, photoelectric noise.

More specifically, in this paper we are interested in problems where \mathbf{H} models space-invariant periodic convolutions in the original image domain. This class of problems are usually termed *image deconvolution* or *image restoration*. The corresponding matrix \mathbf{H} is then square (with $N_x = N_y = N$) block-circulant and can be diagonalized by the 2D discrete Fourier transform (DFT):

$$\mathbf{H} = \mathbf{U}^H \mathbf{D} \mathbf{U}. \quad (2)$$

In the above equation, \mathbf{U} is the matrix that represents the 2D discrete Fourier transform, $\mathbf{U}^H = \mathbf{U}$ is its inverse (since \mathbf{U} is a unitary matrix, that is, $\mathbf{U}\mathbf{U}^H = \mathbf{U}^H\mathbf{U} = \mathbf{I}$, where $(\cdot)^H$ denotes conjugate transpose), and \mathbf{D} is a diagonal matrix containing the DFT coefficients of the convolution operator represented by \mathbf{H} . This means that multiplication by \mathbf{H} can be performed in the discrete Fourier domain with a simple point-wise multiplication (recall that \mathbf{D} is diagonal)

$$\mathbf{H}\mathbf{x} = \mathbf{U}^H \mathbf{D} \mathbf{U}\mathbf{x} = \mathbf{U}^H \mathbf{D} \tilde{\mathbf{x}},$$

where $\tilde{\mathbf{x}} = \mathbf{U}\mathbf{x}$ denotes the DFT of \mathbf{x} .

If matrix \mathbf{H} is not block-circulant, but just block-Toeplitz, it is possible to embed the non-periodic convolution that it represents in a larger periodic convolution and still work in the DFT

domain [16]. Using this technique, all the results and statements made in this paper concerning circulant observation matrices (periodic convolutions) can be extended to the Toeplitz case.

III. REVIEW OF FFT-BASED RECOVERY AND WIENER FILTERING

If \mathbf{H} is invertible (*i.e.*, there are no zeros in the diagonal of \mathbf{D} , thus \mathbf{D}^{-1} exists) we can write $\mathbf{H}^{-1} = \mathbf{U}^H \mathbf{D}^{-1} \mathbf{U}$. Then, if we ignore the presence of noise, we can obtain an estimate of \mathbf{x} as

$$\hat{\mathbf{x}} = \mathbf{U}^H \mathbf{D}^{-1} \mathbf{U} \mathbf{y} = \mathbf{U}^H \mathbf{D}^{-1} \tilde{\mathbf{y}}, \quad (3)$$

where $\tilde{\mathbf{y}} = \mathbf{U} \mathbf{y}$ denotes the DFT of the observation \mathbf{y} . Of course, in practice, the DFT and its inverse are computed by using the *fast Fourier transform* (FFT) algorithm, which requires $O(N \log N)$ operations (where $N = N_x = N_y$ is the number of pixels), and not with matrix multiplications. Consequently, implementing (3) also requires $O(N \log N)$ operations.

In most cases of interest, \mathbf{H} is non-invertible (there are zeros in the diagonal of \mathbf{D}) or at least very ill-conditioned (there are very small values in the diagonal of \mathbf{D}), with direct inversion leading to a severe amplification of the observation noise. Therefore, some regularization procedure is required. A common choice is to adopt a *maximum penalized likelihood estimator* (MPLE)

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} \{ \log p(\mathbf{y}|\mathbf{x}) - \text{pen}(\mathbf{x}) \} = \arg \min_{\mathbf{x}} \{ -\log p(\mathbf{y}|\mathbf{x}) + \text{pen}(\mathbf{x}) \}, \quad (4)$$

where $p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}|\mathbf{H}\mathbf{x}, \sigma^2 \mathbf{I})$ is the likelihood function corresponding to the observation model in (1), and $\text{pen}(\mathbf{x})$ is a penalty function. From a Bayesian perspective, this is a *maximum a posteriori* (MAP) criterion under the prior $p(\mathbf{x})$, such that $\text{pen}(\mathbf{x}) = -\log p(\mathbf{x})$.

If the prior $p(\mathbf{x})$ is Gaussian, with mean $\boldsymbol{\mu}$ (usually zero) and covariance matrix \mathbf{G} , it is well-known (see, for example, [31]) that the MPLE/MAP estimate can be written as

$$\begin{aligned} \hat{\mathbf{x}} &= \arg \min_{\mathbf{x}} \left\{ \frac{1}{\sigma^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|^2 + (\mathbf{x} - \boldsymbol{\mu})^H \mathbf{G}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\} \\ &= \boldsymbol{\mu} + \mathbf{G}\mathbf{H}^H (\sigma^2 \mathbf{I} + \mathbf{H}\mathbf{G}\mathbf{H}^H)^{-1} (\mathbf{y} - \mathbf{H}\boldsymbol{\mu}). \end{aligned} \quad (5)$$

When the covariance of the prior, \mathbf{G} , is also (as the observation matrix \mathbf{H}) block-circulant (meaning that the original image is considered a sample of stationary Gaussian field with periodic

boundary conditions), it is also diagonalized by the DFT and we can write $\mathbf{G} = \mathbf{U}^H \mathbf{C} \mathbf{U}$, where \mathbf{C} is diagonal. In this case, (5) can be implemented in the DFT domain as

$$\hat{\mathbf{x}} = \boldsymbol{\mu} + \mathbf{U}^H \mathbf{C} \mathbf{D}^H (\sigma^2 \mathbf{I} + \mathbf{D} \mathbf{C} \mathbf{D}^H)^{-1} (\mathbf{U} \mathbf{y} - \mathbf{D} \mathbf{U} \boldsymbol{\mu}). \quad (6)$$

Since the matrix being inverted in (6) is diagonal, the leading computational cost is the $O(N \log N)$ corresponding to the FFTs $\mathbf{U} \boldsymbol{\mu}$ and $\mathbf{U} \mathbf{y}$ and to the inverse FFT expressed by the left multiplication by \mathbf{U}^H . Equation (6) is a Wiener filter in the DFT domain [1].

Unfortunately, this FFT-based procedure only discriminates between signal and noise in the frequency domain. It is well-known that real-world images are not well modelled by stationary Gaussian fields. A typical image \mathbf{x} will not admit a sparse Fourier representation; the signal energy may not be concentrated in a small subspace, making it difficult to remove noise and preserve signal simultaneously.

IV. WAVELET-BASED IMAGE RESTORATION

A. Introduction

In wavelet-based estimation, the image \mathbf{x} is re-expressed in terms of an orthogonal wavelet expansion, which typically provides a very sparse representation (a few large coefficients and many very small ones) [23]. Let \mathbf{W} denote the (inverse) discrete wavelet transform (DWT) and let us write $\mathbf{x} = \mathbf{W} \boldsymbol{\theta}$, where $\boldsymbol{\theta}$ is the vector of wavelet coefficients [23]. As above, let us consider an MPLE/MAP criterion, expressed in terms of $\boldsymbol{\theta}$, the wavelet coefficients of the original image, that is, taking the likelihood function to be $p(\mathbf{y}|\boldsymbol{\theta})$. Considering some penalty $pen(\boldsymbol{\theta})$ emphasizing sparsity of the DWT coefficients, the MPLE/MAP estimate is given by

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \{-\log p(\mathbf{y}|\boldsymbol{\theta}) + pen(\boldsymbol{\theta})\} \quad (7)$$

$$= \arg \min_{\boldsymbol{\theta}} \left\{ \frac{\|\mathbf{y} - \mathbf{H} \mathbf{W} \boldsymbol{\theta}\|^2}{2\sigma^2} + pen(\boldsymbol{\theta}) \right\}. \quad (8)$$

The penalty function can be interpreted as minus the logarithm of some (non-Gaussian, sparseness-inducing) prior [26], $pen(\boldsymbol{\theta}) = -\log p(\boldsymbol{\theta})$, as a complexity-based penalty [27], or as a regularization term [2].

When $\mathbf{H} = \mathbf{I}$, that is, for direct denoising problems, wavelet-based methods are extremely efficient, thanks to the fast implementations of the DWT and to the orthogonality of \mathbf{W} (that is, $\mathbf{W}^T \mathbf{W} = \mathbf{W} \mathbf{W}^T = \mathbf{I}$) which allows solving (8) by applying a coefficient-wise denoising rule; moreover, these methods achieve state-of-the-art performance (see [14], [23], [25], [26] and references therein). The very good performance of wavelet-based denoising methods can be traced back to the adequacy of the underlying priors/models of real world images.

Although wavelets have also been shown to be effective in image restoration problems (see [3], [4], [8], [9], [17], [18], [22], [28], [29], [33], [34]), major difficulties arise:

- (a) unlike \mathbf{H} alone, $\mathbf{H}\mathbf{W}$ is not block-circulant, thus it is not diagonalized by the DFT;
- (b) unlike \mathbf{W} alone, $\mathbf{H}\mathbf{W}$ is not orthogonal, thus precluding efficient coefficient-wise rules.

B. Previous Work

In [4], [22], [33], general frameworks aimed at restoration approaches of the form of (8) has been proposed. The results are promising, but the proposed algorithms are very numerically intensive. The iterative method of [28] is also similar in spirit, employing an *ad hoc* thresholding step within an iterative restoration algorithm. In certain exceptional cases in which the operator \mathbf{H} is scale-homogeneous, and hence (approximately) diagonalized by \mathbf{W} , the so-called *wavelet-vaguelette* procedure developed by Donoho [9] leads to very efficient threshold restoration procedures. However, most convolution operators are not scale-invariant and thus the wavelet-vaguelette procedure is not applicable.

An adaptation of the *wavelet-vaguelette* approach, based on wavelet-packets designed to match the frequency behavior of certain convolutions, was proposed in [18]. This method was extended to a complex wavelet *hidden Markov tree* (see [6]) scheme in [17]. Although these methods are computationally fast, they are not applicable to most convolutions and, moreover, choosing the (image) basis to conform to the operator is exactly what wavelet methods set out to avoid in the first place. The wavelet packets matched to the frequency behavior of the convolution operator may not match image structure as well as a conventional wavelet basis.

Other methods for more general deconvolution problems have been proposed. In [3], the approach is to adapt the linear filtering spatially, based on an edge detection test. The algorithm presented in [29] combines Fourier domain regularization with wavelet domain thresholding, in a non-iterative fashion, with very good results. Recently, an iterative method using preconditioned conjugate gradient was proposed in [8]; the method achieves very good results, but it requires complex wavelet transforms and a complicated initialization procedure based on another wavelet-based restoration method (namely the one proposed in [29]).

Finally, we mention that EM and EM-type algorithms have been previously used in image restoration and reconstruction, with non-wavelet-based formulations (*e.g.*, [11], [12], [19]).

V. THE BEST OF BOTH WORLDS

The approach proposed in this paper is able to use the best of the wavelet and Fourier worlds in image deconvolution problems. The speed and convenience of the FFT-based Wiener filter, which is well matched to the observation model, and the adequacy of wavelet-based image models.

A. An Equivalent Model and the EM Algorithm

Let us write the observation model in (1) with respect to the DWT coefficients $\boldsymbol{\theta}$ (recall that $\mathbf{x} = \mathbf{W}\boldsymbol{\theta}$):

$$\mathbf{y} = \mathbf{H}\mathbf{W}\boldsymbol{\theta} + \mathbf{n}. \quad (9)$$

As mentioned above, this equation clearly shows where the difficulties come from: although \mathbf{H} is diagonalized by the DFT, $\mathbf{H}\mathbf{W}$ is not, and so FFT-based methods are not directly applicable. To overcome this problem, we propose decomposing the white Gaussian noise \mathbf{n} into the sum of two different Gaussian noises (one of which is non-white), *i.e.*,

$$\mathbf{n} = \alpha \mathbf{H} \mathbf{n}_1 + \mathbf{n}_2, \quad (10)$$

where α is a positive parameter, and \mathbf{n}_1 and \mathbf{n}_2 are independent noises such that

$$\begin{aligned} p(\mathbf{n}_1) &= \mathcal{N}(\mathbf{n}_1|0, \mathbf{I}) \\ p(\mathbf{n}_2) &= \mathcal{N}(\mathbf{n}_2|0, \sigma^2\mathbf{I} - \alpha^2\mathbf{H}\mathbf{H}^T). \end{aligned}$$

Notice that the covariance of $\alpha \mathbf{H} \mathbf{n}_1 + \mathbf{n}_2$ is $\alpha^2 \mathbf{H} \mathbf{H}^T + \sigma^2 \mathbf{I} - \alpha^2 \mathbf{H} \mathbf{H}^T = \sigma^2 \mathbf{I}$, as required. For $(\sigma^2 \mathbf{I} - \alpha^2 \mathbf{H} \mathbf{H}^T)$ to be semi-positive definite (thus a valid covariance matrix), we must have $\alpha^2 \leq \sigma^2 / \lambda_1$, where λ_1 is the largest eigenvalue of $\mathbf{H} \mathbf{H}^T$. With a normalized (total mass equal to one) and periodic (corresponding to a block-circulant¹ \mathbf{H}) blur, we have $\lambda_1 = 1$, and the condition simplifies to $\alpha^2 \leq \sigma^2$. The idea behind the proposed noise decomposition is that it allows the introduction of a hidden image \mathbf{z} which decouples the denoising from the deconvolution, as next described. Notice that using \mathbf{n}_1 and \mathbf{n}_2 , we can decompose the observation model as

$$\begin{cases} \mathbf{z} &= \mathbf{W} \boldsymbol{\theta} + \alpha \mathbf{n}_1 \\ \mathbf{y} &= \mathbf{H} \mathbf{z} + \mathbf{n}_2. \end{cases} \quad (11)$$

Clearly, if we had \mathbf{z} , we would have a pure denoising problem with white noise (the first equation in (11)). This observation is the key to our approach, since it suggests treating \mathbf{z} as missing data and estimating $\boldsymbol{\theta}$ via the EM algorithm (see, *e.g.*, [7], [24]). Recall that the EM algorithm is a means of obtaining MAP/MPLE estimates (of which maximum likelihood is a particular case) of a parameter (see (7)) in cases where the penalized log-likelihood $\log p(\mathbf{y} | \boldsymbol{\theta}) - \text{pen}(\boldsymbol{\theta})$ is hard to maximize, but the so-called *complete penalized log-likelihood* $\log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) - \text{pen}(\boldsymbol{\theta})$ would be easy to maximize if we had \mathbf{z} . The EM algorithm produces a sequence of estimates $\{\hat{\boldsymbol{\theta}}^{(t)}, t = 0, 1, 2, \dots\}$ by alternating two steps (until some stopping criterion is met):

- **E-step:** Computes the conditional expectation of the log-likelihood of the complete data, given the observed data and the current estimate $\hat{\boldsymbol{\theta}}^{(t)}$. The result is the so-called *Q-function*:

$$Q(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(t)}) \equiv E \left[\log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) \mid \mathbf{y}, \hat{\boldsymbol{\theta}}^{(t)} \right]. \quad (12)$$

- **M-step:** Updates the estimate according to

$$\hat{\boldsymbol{\theta}}^{(t+1)} = \arg \min_{\boldsymbol{\theta}} \{-Q(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(t)}) + \text{pen}(\boldsymbol{\theta})\}. \quad (13)$$

¹If \mathbf{H} is not block-circulant, but block-Toeplitz, as long as the blur satisfies some very mild conditions, the eigenvalues are, asymptotically (in the size of the matrix) the same (see [15] and references therein); with blurs that are much smaller than the image, the eigenvalues of the corresponding Toeplitz or circulant matrices are then roughly the same.

It is well known (see, *e.g.*, [7], [24]) that each iteration of the EM algorithm is guaranteed to increase the penalized log-likelihood, that is,

$$-\log p(\mathbf{y}|\hat{\boldsymbol{\theta}}^{(t+1)}) + \text{pen}(\hat{\boldsymbol{\theta}}^{(t+1)}) \leq -\log p(\mathbf{y}|\hat{\boldsymbol{\theta}}^{(t)}) + \text{pen}(\hat{\boldsymbol{\theta}}^{(t)}).$$

Next, we derive the specific formulas for the E and M steps, for our deconvolution problem.

B. The E-Step: FFT-Based Estimation

The complete-data likelihood is $p(\mathbf{y}, \mathbf{z}|\boldsymbol{\theta}) = p(\mathbf{y}|\mathbf{z}, \boldsymbol{\theta})p(\mathbf{z}|\boldsymbol{\theta}) = p(\mathbf{y}|\mathbf{z})p(\mathbf{z}|\boldsymbol{\theta})$, because, conditioned on \mathbf{z} , \mathbf{y} is independent of $\boldsymbol{\theta}$ (see (11)). Since $\mathbf{z} = \mathbf{W}\boldsymbol{\theta} + \alpha \mathbf{n}_1$, where $\alpha \mathbf{n}_1$ is zero-mean with covariance $\alpha^2 \mathbf{I}$, we simply have

$$\begin{aligned} \log p(\mathbf{y}, \mathbf{z}|\boldsymbol{\theta}) &= -\frac{\|\mathbf{W}\boldsymbol{\theta} - \mathbf{z}\|^2}{2\alpha^2} + K_1 \\ &= -\frac{\boldsymbol{\theta}^T \mathbf{W}^T \mathbf{W} \boldsymbol{\theta} - 2\boldsymbol{\theta}^T \mathbf{W}^T \mathbf{z}}{2\alpha^2} + K_2, \end{aligned} \quad (14)$$

where K_1 and K_2 are constants that do not depend on $\boldsymbol{\theta}$. This shows that the complete-data log-likelihood is linear with respect to the missing data \mathbf{z} . Consequently, all that is required in the E-step is to compute the conditional expectation of \mathbf{z} , given the observed data \mathbf{y} and current parameter estimate $\hat{\boldsymbol{\theta}}^{(t)}$,

$$\hat{\mathbf{z}}^{(t)} \equiv E[\mathbf{z}|\mathbf{y}, \hat{\boldsymbol{\theta}}^{(t)}] = \int \mathbf{z} p(\mathbf{z}|\mathbf{y}, \hat{\boldsymbol{\theta}}^{(t)}) d\mathbf{z}, \quad (15)$$

and plug it into the complete-data log-likelihood to obtain

$$Q(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(t)}) = -\frac{\boldsymbol{\theta}^T \mathbf{W}^T \mathbf{W} \boldsymbol{\theta} - 2\boldsymbol{\theta}^T \mathbf{W}^T \hat{\mathbf{z}}^{(t)}}{2\alpha^2} + K_2 = -\frac{\|\mathbf{W}\boldsymbol{\theta} - \hat{\mathbf{z}}^{(t)}\|^2}{2\alpha^2} + K_1. \quad (16)$$

Since $p(\mathbf{y}|\mathbf{z}) = \mathcal{N}(\mathbf{y}|\mathbf{H}\mathbf{z}, \sigma^2 \mathbf{I} - \alpha^2 \mathbf{H}\mathbf{H}^T)$ and $p(\mathbf{z}|\hat{\boldsymbol{\theta}}^{(t)}) = \mathcal{N}(\mathbf{z}|\mathbf{W}\hat{\boldsymbol{\theta}}^{(t)}, \alpha^2 \mathbf{I})$, then $p(\mathbf{z}|\mathbf{y}, \hat{\boldsymbol{\theta}}^{(t)}) \propto p(\mathbf{y}|\mathbf{z})p(\mathbf{z}|\hat{\boldsymbol{\theta}}^{(t)})$ is also Gaussian, with mean given by (see, *e.g.*, [31])

$$\hat{\mathbf{z}}^{(t)} = \mathbf{W}\hat{\boldsymbol{\theta}}^{(t)} + \frac{\alpha^2}{\sigma^2} \mathbf{H}^T (\mathbf{y} - \mathbf{H}\mathbf{W}\hat{\boldsymbol{\theta}}^{(t)}) \quad (17)$$

$$= \mathbf{W}\hat{\boldsymbol{\theta}}^{(t)} + \frac{\alpha^2}{\sigma^2} \mathbf{U}^H \mathbf{D}^H (\mathbf{U}\mathbf{y} - \mathbf{D}\mathbf{U}\mathbf{W}\hat{\boldsymbol{\theta}}^{(t)}), \quad (18)$$

which can be efficiently implemented by FFT (recall that $\mathbf{U}^H \mathbf{D}^H \mathbf{U} = \mathbf{H}^T$ and $\mathbf{U}^H \mathbf{D}^H \mathbf{D}\mathbf{U} = \mathbf{H}^T \mathbf{H}$). Notice that since $\hat{\mathbf{x}}^{(t)} \equiv \mathbf{W}\hat{\boldsymbol{\theta}}^{(t)}$ can be seen as the current estimate of the true image \mathbf{x} ,

we can write the E-step as

$$\hat{\mathbf{z}}^{(t)} = \hat{\mathbf{x}}^{(t)} + \frac{\alpha^2}{\sigma^2} \mathbf{H}^T (\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}^{(t)}), \quad (19)$$

revealing its similarity with a Landweber iteration for solving $\mathbf{H}\mathbf{x} = \mathbf{y}$ [20], [32]. Of course this is just the E-step; the complete EM algorithm is not a Landweber algorithm.

C. M-Step: Wavelet-Based Denoising

In the M-step, the parameter estimate is updated as shown in (13), where $Q(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(t)})$ is as given by (16) with $\hat{\mathbf{z}}^{(t)}$ computed according to (18):

$$\hat{\boldsymbol{\theta}}^{(t+1)} = \arg \min_{\boldsymbol{\theta}} \left\{ \frac{\|\mathbf{W}\boldsymbol{\theta} - \hat{\mathbf{z}}^{(t)}\|^2}{2\alpha^2} + pen(\boldsymbol{\theta}) \right\}. \quad (20)$$

This is simply a MPLE/MAP estimate of $\boldsymbol{\theta}$, under the prior $p(\boldsymbol{\theta})$, for a “direct” observation denoising problem: we observe $\hat{\mathbf{z}}^{(t)} \sim \mathcal{N}(\mathbf{W}\boldsymbol{\theta}, \alpha^2\mathbf{I})$. Because the wavelet transform is orthogonal we have $\|\mathbf{W}\boldsymbol{\theta} - \hat{\mathbf{z}}^{(t)}\|^2 = \|\boldsymbol{\theta} - \hat{\boldsymbol{\omega}}^{(t)}\|^2$, where $\hat{\boldsymbol{\omega}}^{(t)} \equiv \mathbf{W}^T \hat{\mathbf{z}}^{(t)}$ denotes the DWT transform of $\hat{\mathbf{z}}^{(t)}$. Thus, the M-Step can be computed by applying the corresponding denoising rule to $\hat{\boldsymbol{\omega}}^{(t)}$:

$$\hat{\boldsymbol{\theta}}^{(t+1)} = \arg \min_{\boldsymbol{\theta}} \left\{ \|\boldsymbol{\theta} - \hat{\boldsymbol{\omega}}^{(t)}\|^2 + 2\alpha^2 pen(\boldsymbol{\theta}) \right\}. \quad (21)$$

For example, under a i.i.d. Laplacian prior on the wavelet coefficients,

$$p(\boldsymbol{\theta}) \propto \exp\{-\tau\|\boldsymbol{\theta}\|_1\} \Rightarrow pen(\boldsymbol{\theta}) = \tau \sum_i |\theta_i|, \quad (22)$$

(where $\|\boldsymbol{\theta}\|_1 = \sum_i |\theta_i|$ denotes the l_1 norm), $\hat{\boldsymbol{\theta}}^{(t+1)}$ is obtained by applying a *soft-threshold* function to $\hat{\boldsymbol{\omega}}^{(t)}$, the wavelet coefficients of $\hat{\mathbf{z}}^{(t)}$ [26]. More specifically, each component of $\hat{\boldsymbol{\theta}}^{(t+1)}$ is obtained separately according to

$$\hat{\theta}_i^{(t+1)} = \text{sgn}\left(\hat{\omega}_i^{(t)}\right) \left(|\hat{\omega}_i^{(t)}| - \tau\alpha^2\right)_+ \quad (23)$$

where $(\cdot)_+$ denotes the *positive part operator*, defined as $(x)_+ = \max\{x, 0\}$, and $\text{sgn}(\cdot)$ is the *sign function*, defined as $\text{sgn}(x) = 1$, if $x > 0$, and $\text{sgn}(x) = -1$, if $x < 0$. Other priors or complexity penalties will lead to different wavelet denoising rules in the M-Step [14], [23], [26], [27].

D. Computational Complexity

The computational complexity of the M-Step is dominated by the DWT, usually $O(N)$ for an orthogonal DWT. The computational load of the E-step is dominated by the $O(N \log N)$ cost of the FFT. The cost of each iteration of the complete EM algorithm is thus $O(N \log N)$.

E. Some Comments

A very important feature of this EM algorithm is that any wavelet denoising procedure that can be interpreted as an MPLE/MAP rule can be employed in the M-Step. For example, $p(\boldsymbol{\theta})$ could correspond to a hidden Markov tree model [6] or to a locally adaptive model [25]; however, in those cases, the M-step would not be as simple as a fixed nonlinear thresholding rule. We can also use the denoising rule that we have proposed in [13], [14], since although it was originally derived from an empirical-Bayes approach, we have shown that it corresponds to an MPLE/MAP estimate under a prior of a particular form [14]; the rule is given by

$$\hat{\theta}_i^{(t+1)} = \frac{\left((\hat{\omega}_i^{(t)})^2 - 3\alpha^2 \right)_+}{\hat{\omega}_i^{(t)}}. \quad (24)$$

Let \mathcal{D} denote whichever denoising operation is applied to the wavelet coefficients (such as (23) or (24)), and \mathcal{P} the resulting denoising procedure applied to some image \mathbf{v} , that is,

$$\mathcal{P}(\mathbf{v}) \equiv \mathbf{W}\mathcal{D}(\mathbf{W}^T \mathbf{v}). \quad (25)$$

With this notation, we can write compact a expression for each iteration of the EM algorithm

$$\hat{\mathbf{x}}^{(t+1)} = \mathcal{P} \left(\hat{\mathbf{x}}^{(t)} + \frac{\alpha^2}{\sigma^2} \mathbf{H}^T (\mathbf{y} - \mathbf{H}\hat{\mathbf{x}}^{(t)}) \right), \quad (26)$$

which can be interpreted as a Landweber iteration followed by a wavelet-based denoising step.

Of course the choice of α affects the rate of convergence of the algorithm. The standard theory of the rate of convergence of EM, based on the information matrices (see [24]), suggests that α should be made as large as possible. Since we must have $\alpha^2 \leq \sigma^2$ to have a meaningful EM algorithm (see Subsection V-A), a reasonable choice is $\alpha^2 = \sigma^2$. Although the analysis of the rate of convergence based on the information matrices can only be performed ignoring the penalty

terms, since these may be non-differentiable, we found experimentally that $\alpha^2 = \sigma^2$ is indeed a good choice. Another line of thought uses the similarity of (26) with a Landweber iteration; ignoring the denoising operator \mathcal{P} suggests that we have a convergent algorithm for $\alpha^2/\sigma^2 \leq 2$ (rather than the 1 that results from the EM interpretation).

Finally, let us summarize the several very attractive features of this approach:

- the computational complexity of each iteration is $O(N \log N)$;
- we can employ *any* orthogonal wavelet basis;
- we can employ *any* wavelet-based penalization.

VI. EXTENSION TO UNKNOWN NOISE VARIANCE

Up to this point, we have assumed that the noise variance σ^2 is known in advance. We now present an extension of the proposed algorithm which also estimates σ^2 . This is simply done by inserting an additional step in which the noise variance estimate is updated based on the current estimate of the true image $\hat{\mathbf{x}}(t) \equiv \mathbf{W} \hat{\boldsymbol{\theta}}(t)$. The complete algorithm is now defined by two steps:

- **EM step:** equation (26);
- **Noise variance update:**

$$\widehat{\sigma^2}^{(t+1)} = \frac{\|\mathbf{H} \hat{\mathbf{x}}^{(t+1)} - \mathbf{y}\|^2}{N}. \quad (27)$$

The complete algorithm is not an EM algorithm, but it is also guaranteed to increase the penalized likelihood function. To see that this is true, let us denote the penalized negative log-likelihood being minimized (which is now also a function of σ^2) as

$$\mathcal{L}(\boldsymbol{\theta}, \sigma^2) = \frac{N}{2} \log \sigma^2 + \frac{\|\mathbf{H} \mathbf{W} \boldsymbol{\theta} - \mathbf{y}\|^2}{2\sigma^2} + \text{pen}(\boldsymbol{\theta}). \quad (28)$$

Concerning the EM step, we know that $\mathcal{L}(\hat{\boldsymbol{\theta}}^{(t+1)}, \widehat{\sigma^2}^{(t)}) \leq \mathcal{L}(\hat{\boldsymbol{\theta}}^{(t)}, \widehat{\sigma^2}^{(t)})$, due to the monotonicity properties of the EM algorithm [24]. The noise variance updating step is simply a maximum likelihood estimate of σ^2 , with the estimate of $\boldsymbol{\theta}$ fixed at $\hat{\boldsymbol{\theta}}^{(t+1)}$,

$$\widehat{\sigma^2}^{(t+1)} = \frac{\|\mathbf{H} \mathbf{W} \hat{\boldsymbol{\theta}}^{(t+1)} - \mathbf{y}\|^2}{N} = \arg \min_{\sigma^2} \mathcal{L}(\hat{\boldsymbol{\theta}}^{(t+1)}, \sigma^2),$$

since $pen(\boldsymbol{\theta})$ does not depend on σ^2 . Accordingly, we have $\mathcal{L}(\widehat{\boldsymbol{\theta}}^{(t+1)}, \widehat{\sigma}^{2(t+1)}) \leq \mathcal{L}(\widehat{\boldsymbol{\theta}}^{(t+1)}, \widehat{\sigma}^{2(t)})$. In conclusion, since both steps are guaranteed not to decrease the penalized log-likelihood function, so is their combination.

VII. CONVERGENCE ANALYSIS OF THE EM ALGORITHM

A general, basic property of an EM algorithm is that it generates a sequence of non-decreasing (penalized) likelihood values [24]. EM iterations produce a sequence of images, each of which has a penalized likelihood value greater than or equal to that of the preceding image. This is a very desirable property, but several questions remain. (1) Does the sequence (of penalized likelihood values) converge to the maximum of the penalized likelihood function? (2) Does the corresponding sequence of images converge to a fixed image and is this limit (assuming it exists) unique? This section explores these issues. First, we consider the conditions under which the EM algorithm converges to a stationary point of the penalized likelihood function. Second, we investigate the convexity of the penalized negative log-likelihood function and establish conditions under which the EM algorithm converges to a unique solution.

A. Convergence to a Stationary Point

The results in [35] guarantee that the EM algorithm converges to a stationary point (local maximum or saddle-point) of the penalized likelihood function under fairly mild conditions. Theorem 2 of [35] shows that all limit points of the EM algorithm are stationary points of the penalized likelihood function, provided that $Q(\boldsymbol{\theta}, \widehat{\boldsymbol{\theta}}^{(t)})$ and $pen(\boldsymbol{\theta})$ are continuous in both $\boldsymbol{\theta}$ and $\widehat{\boldsymbol{\theta}}$. This condition is clearly met by the expected complete-data log-likelihood $Q(\boldsymbol{\theta}, \widehat{\boldsymbol{\theta}}^{(t)})$. The penalty function $pen(\boldsymbol{\theta})$ also needs to be continuous in order to guarantee convergence to a stationary point. This precludes the use of the conventional hard-threshold function, but both the soft-threshold rule (23) and our rule in (24) correspond to continuous penalty functions (log-priors). To summarize, if the penalty function underlying the nonlinear shrinkage/threshold function employed in the M-Step is continuous in $\boldsymbol{\theta}$, then the EM algorithm converges to a stationary

point of the penalized log-likelihood. The limit points may be local maxima or saddle-points; it is difficult to guarantee convergence to a local maximum without further assumptions. Such conditions are investigated next.

B. Convergence to a Global Maximum

Let us begin by considering the case in which \mathbf{H} is invertible. Under this assumption, the negative log-likelihood term of (8) is strictly convex in $\boldsymbol{\theta}$. Now if the penalty function is also convex (not necessarily strictly so), then the penalized negative log-likelihood function is strictly convex in $\boldsymbol{\theta}$. For example, the log-Laplacian penalty function, leading to the soft-threshold rule, is convex (not strictly) in $\boldsymbol{\theta}$. Strict convexity of the penalized negative log-likelihood function implies that there is only one stationary point, the global maximum. Thus, under the continuity conditions discussed above, the EM algorithm is guaranteed to converge to the global maximum. Note that the uniqueness of the maximum point guarantees that the sequence of images produced by the EM algorithm converges to the global maximum penalized likelihood image restoration.

Next consider situations when \mathbf{H} is not invertible. For examples, \mathbf{H} is not invertible if the DFT of the underlying point spread response is zero at some point(s). In such cases, the negative log-likelihood term of (8) is convex, but not strictly, in $\boldsymbol{\theta}$. If the penalty function is also convex (but not strictly so), then the sequences of penalized log-likelihood values produced by the EM algorithms will converge to their respective global maximum penalized log-likelihood values. This follows from the EM convergence results of Wu [35], since all stationary points of a convex function are global minima. However, since there may be many global minima, the EM algorithms may not converge to fixed images (they are only guaranteed to converge to their respective sets of images corresponding to global minima). If it does converge to a fixed image (this limit could depend on the initialization of the algorithm), then that image maximizes the penalized likelihood criterion.

If the penalty function is strictly convex, then the EM algorithm is guaranteed to converge to the unique maximum penalized likelihood value and a unique optimal image. This also follows

from the EM convergence results [35]; the unique stationary point of a strictly convex function is the global minimum. So far, the only convex penalty function we have considered is the log-Laplacian (leading to the soft-threshold rule), but even this penalty function is not strictly convex, since its growth with the absolute value of the argument is linear. The following modification of the log-Laplacian leads to a strictly convex penalty function and a threshold rule nearly the same as the soft-threshold function, except that it is differentiable at all points. Instead of the log-Laplacian penalty, which has the form $-\log e^{-\alpha|\theta|} = \alpha|\theta|$, consider

$$pen(\theta) = -\log e^{-\eta\sqrt{\theta^2+\beta^2}} = \eta\sqrt{\theta^2+\beta^2}, \quad (29)$$

for some small number β . Notice that as $\beta \rightarrow 0$, this penalty tends to the log-Laplacian. However, for every $\alpha, \beta > 0$ this penalty is strictly convex, since $d^2(\eta\sqrt{\theta^2+\beta^2})/d\theta^2 > 0$. The difference between the threshold rule induced by the penalty (29) and the soft-threshold is that the former makes a smooth transition across the threshold level, as shown in Figure 1.




Fig. 1. Soft-threshold function (dashed) and modified soft-threshold function (solid) with threshold level set at 1 and $\beta = 1$. If $\beta = 0.1$, then the difference between the soft-threshold function and the modified soft-threshold function are indistinguishable to the naked eye at this scale.

C. Summary of Convergence Results

The following four points summarize the convergence properties of our EM algorithm.

1. If the penalty $pen(\boldsymbol{\theta})$ is a continuous function of $\boldsymbol{\theta}$, then each iteration of the EM algorithm produces an image with a penalized likelihood value greater than or equal to the previous image.

2. If the penalty function is also convex (but not strictly convex) in θ , then the sequence of penalized log likelihood values converges to the global maximum. However, since there may be many global maxima, the EM algorithm may not converge to a fixed image. If it does converge to a fixed image, then that image maximizes the penalized likelihood criterion.
3. The EM algorithm converges to the unique, globally optimal solution (fixed image) of the penalized likelihood criterion if either of the conditions below are met:
 - a) \mathbf{H} is invertible and the penalty function is convex (*e.g.*, soft-threshold);
 - b) the penalty function is strictly convex (*e.g.*, the modified soft-threshold penalty (29)).
4. Recall that the EM algorithm coupled with the adaptive updates of the noise variance, given by (27), produce non-decreasing sequences of penalized likelihood values (with the noise variance σ^2 treated as an unknown parameter to be inferred jointly with θ). However, the corresponding penalized negative log likelihood function is non-convex and convergence is no longer guaranteed in this case.

VIII. EXTENSION TO TRANSLATION-INVARIANT RESTORATION

It is well known that the dyadic image partitioning underlying the orthogonal DWT can cause blocky artifacts in the processed images. In denoising problems, translation-invariant approaches have been shown to significantly reduce these artifacts and are routinely used instead of the orthogonal DWT [5], [14], [21]. The standard way to achieve translation invariance in denoising is to use a redundant transform, called the translation-invariant DWT (TI-DWT), which corresponds to computing the inner products between the image and all (circularly) translated versions of the wavelet basis functions. Denoising is accomplished by thresholding as usual and then averaging the results. Working with all possible shifts of the discrete wavelet basis functions, rather than the dyadic shifts underlying the orthogonal DWT basis functions, helps to reduce blocky artifacts and achieves better denoising performance [5], [14], [21].

In this paper, we consider three ways to achieve translation invariance in our iterative image deconvolution algorithm, which we describe in the following three subsections.

A. Translation Invariance via Undecimated DWT

The TI-DWT is an over-complete transform based on N orthogonal DWTs. Each of the N DWTs is comprised of circularly shifted versions of the discrete DWT basis functions. Let \mathbf{W}_0^T be an orthogonal DWT matrix. Let $i \in \{0, \dots, N-1\}$ index all possible circular image shifts; let \mathbf{W}_i^T denote a DWT matrix with the i -th shift applied to all the basis functions in \mathbf{W}_0^T . With this notation, the TI-DWT matrix can be written as

$$\mathbf{W}^T = \frac{1}{\sqrt{N}} [\mathbf{W}_0 \cdots \mathbf{W}_{N-1}]^T. \quad (30)$$

Since the TI-DWT is not invertible, the pseudo-inverse

$$\mathbf{W} = \frac{1}{\sqrt{N}} [\mathbf{W}_0 \cdots \mathbf{W}_{N-1}] \quad (31)$$

is standardly used to transform the redundant set of coefficients back to the image space. Notice that if \mathbf{x} denotes any image, then

$$\mathbf{W}\mathbf{W}^T\mathbf{x} = \frac{1}{N} [\mathbf{W}_0 \cdots \mathbf{W}_{N-1}] \begin{bmatrix} \mathbf{W}_0^T \\ \vdots \\ \mathbf{W}_{N-1}^T \end{bmatrix} \mathbf{x} = \frac{1}{N} \sum_{i=0}^{N-1} \mathbf{W}_i \mathbf{W}_i^T \mathbf{x} = \mathbf{x},$$

because $\mathbf{W}_i \mathbf{W}_i^T = \mathbf{I}$, thus $\mathbf{W}\mathbf{W}^T = \mathbf{I}$. However, $\mathbf{W}^T\mathbf{W} \neq \mathbf{I}$ and thus \mathbf{W} is not orthogonal.

When \mathbf{W} corresponds to a TI-DWT², the M-Step of our EM algorithm can not be simplified as in (21). However, as is common in denoising [5], [14], [21], we can ignore this fact and still use (21) as if \mathbf{W} were orthogonal. The resulting method is no longer an EM algorithm but, as will be shown below, it leads to excellent image restoration results.

Recall that the coefficients of the TI-DWT can be efficiently computed using the so-called undecimated DWT (UDWT), which simply eliminates the down-sampling process in the filter-bank implementation of a wavelet transform [21]. The TI-DWT produces N^2 coefficients in total, but only $N \log N$ values are unique because certain shifts generate the same inner products between

²A similar complication arises if the orthogonal DWT is replaced by a biorthogonal DWT, but we will not investigate that problem here.

the image and basis functions. The filter-bank implementation of the UDWT produces only the $N \log N$ unique coefficients, and requires $O(N \log N)$ operations. Thus, the computational complexity of each partial optimization in the M-Step is $O(N \log N)$.

Summarizing, our first approach to TI restoration consists simply in keeping the same E-step and M-step, but letting $\hat{\omega}^{(t)}$ in (21) be the UDWT of $\hat{\mathbf{z}}^{(t)}$, rather than its orthogonal DWT.

B. Translation Invariance via Random Shifts

Another possible way to achieve some level of translation invariance consists in choosing a randomly shifted DWT at each iteration. Formally, at each M-step, we let i be a randomly chosen circular shift. Then, we compute an i -shifted orthogonal DWT of $\hat{\mathbf{z}}^{(t)}$, that is $\hat{\omega}^{(t)} \equiv \mathbf{W}_i^T \hat{\mathbf{z}}^{(t)}$ and apply the original denoising step (21). With respect to the UDWT-based approach described in the previous subsection, this method has the advantage of employing an orthogonal DWT, which has $O(N)$ computational cost, rather than the $O(N \log N)$ cost associated with the UDWT.

This method is of course not an EM algorithm. Although the M-step is exact, it corresponds to using a different penalty/prior at each iteration; accordingly, the resulting algorithm can not be interpreted as maximizing some penalized likelihood (or a posteriori probability function).

As shown by the experiments reported below, this method almost always leads to results very close to those obtained by the UDWT-based method.

C. Translation Invariance via a Generalized EM Algorithm

Although both TI restoration methods described above perform well, none of the two is a true EM algorithm, thus they don't have any monotonicity or convergence guarantee. Our third approach to TI restoration consists in using the UDWT but, rather than keeping the original form of the M-step, we change it to recover the monotonicity properties of the algorithm. Specifically, we derive a so-called *generalized EM* (GEM) algorithm, in which the exact maximization performed in the M-step is replaced by a weaker condition:

$$L(\hat{\boldsymbol{\theta}}^{(t+1)}, \hat{\boldsymbol{\theta}}^{(t)}) \leq L(\hat{\boldsymbol{\theta}}^{(t)}, \hat{\boldsymbol{\theta}}^{(t)}) \quad (32)$$

where

$$L(\boldsymbol{\theta}, \widehat{\boldsymbol{\theta}}^{(t)}) \equiv -2\alpha^2 \left(Q(\boldsymbol{\theta}, \widehat{\boldsymbol{\theta}}^{(t)}) + \text{pen}(\boldsymbol{\theta}) \right) = \|\widehat{\mathbf{z}}^{(t)} - \mathbf{W}\boldsymbol{\theta}\|^2 + 2\alpha^2 \text{pen}(\boldsymbol{\theta}) \quad (33)$$

is the function to be minimized in the M-Step (see (20)), and $\boldsymbol{\theta}$ denotes the vector of $N \log N$ unique coefficients associated with the UDWT. The GEM algorithm possesses the same basic monotonicity and convergence properties as the the standard EM [24], [35].

As above, \mathbf{W}_l denotes the orthogonal inverse DWT matrix at an arbitrary shift l , and $\boldsymbol{\theta}_l$ denotes the corresponding set of N coefficients. Writing $\boldsymbol{\theta} = (\boldsymbol{\theta}_l, \boldsymbol{\theta}_{\bar{l}})$, where $\boldsymbol{\theta}_{\bar{l}}$ are the $N(\log N - 1)$ coefficients not associated with the basis functions in \mathbf{W}_l , we have

$$\mathbf{W}\boldsymbol{\theta} = \mathbf{W}_l \boldsymbol{\theta}_l + \mathbf{W}_{\bar{l}} \boldsymbol{\theta}_{\bar{l}}, \quad (34)$$

where $\mathbf{W}_{\bar{l}}$ is composed of the basis functions not in \mathbf{W}_l .

The generalized M-step is obtained by optimizing $L(\boldsymbol{\theta}, \widehat{\boldsymbol{\theta}}^{(t)})$ with respect to $\boldsymbol{\theta}_l$ alone, keeping $\boldsymbol{\theta}_{\bar{l}}$ fixed. To this end, notice that we can write (33) as

$$L(\boldsymbol{\theta}, \widehat{\boldsymbol{\theta}}^{(t)}) = \|\mathbf{W}_l \boldsymbol{\theta}_l + \mathbf{W}_{\bar{l}} \boldsymbol{\theta}_{\bar{l}} - \widehat{\mathbf{z}}^{(t)}\|^2 + 2\alpha^2 \text{pen}((\boldsymbol{\theta}_l, \boldsymbol{\theta}_{\bar{l}})) \quad (35)$$

$$= \|\mathbf{W}_l \boldsymbol{\theta}_l - \mathbf{e}^{(t)}\|^2 + 2\alpha^2 \text{pen}(\boldsymbol{\theta}_l) + 2\alpha^2 \text{pen}(\boldsymbol{\theta}_{\bar{l}}), \quad (36)$$

where $\mathbf{e}^{(t)} = \widehat{\mathbf{z}}^{(t)} - \mathbf{W}_{\bar{l}} \boldsymbol{\theta}_{\bar{l}}^{(t)}$, and where we are assuming a separable penalty function. Then, the generalized M-step is performed by choosing some $l \in \{0, 1, \dots, N-1\}$, either randomly or following some prespecified schedule and then letting

$$\widehat{\boldsymbol{\theta}}_{\bar{l}}^{(t+1)} = \widehat{\boldsymbol{\theta}}_{\bar{l}}^{(t)} \quad (37)$$

$$\widehat{\boldsymbol{\theta}}_l^{(t+1)} = \arg \min_{\boldsymbol{\theta}_l} \left\{ \|\mathbf{W}_l \boldsymbol{\theta}_l - \mathbf{e}^{(t)}\|^2 + 2\alpha^2 \text{pen}(\boldsymbol{\theta}_l) \right\}; \quad (38)$$

finally, we set $\widehat{\boldsymbol{\theta}}^{(t+1)} = (\widehat{\boldsymbol{\theta}}_l^{(t+1)}, \widehat{\boldsymbol{\theta}}_{\bar{l}}^{(t+1)})$. This $\widehat{\boldsymbol{\theta}}^{(t+1)}$ does verify the GEM condition (32):

$$L(\widehat{\boldsymbol{\theta}}^{(t+1)}, \widehat{\boldsymbol{\theta}}^{(t)}) = \|\mathbf{W}_l \widehat{\boldsymbol{\theta}}_l^{(t+1)} - \mathbf{e}^{(t)}\|^2 + 2\alpha^2 \text{pen}(\widehat{\boldsymbol{\theta}}_l^{(t+1)}) + 2\alpha^2 \text{pen}(\widehat{\boldsymbol{\theta}}_{\bar{l}}^{(t)}) \quad (39)$$

$$= \min_{\boldsymbol{\theta}_l} \left\{ \|\mathbf{W}_l \boldsymbol{\theta}_l - \mathbf{e}^{(t)}\|^2 + 2\alpha^2 \text{pen}(\boldsymbol{\theta}_l) \right\} + 2\alpha^2 \text{pen}(\widehat{\boldsymbol{\theta}}_{\bar{l}}^{(t)}) \quad (40)$$

$$\leq \|\mathbf{W}_l \widehat{\boldsymbol{\theta}}_l^{(t)} - \mathbf{e}^{(t)}\|^2 + 2\alpha^2 \text{pen}(\widehat{\boldsymbol{\theta}}_l^{(t)}) + 2\alpha^2 \text{pen}(\widehat{\boldsymbol{\theta}}_{\bar{l}}^{(t)}) \quad (41)$$

$$= L(\widehat{\boldsymbol{\theta}}^{(t)}, \widehat{\boldsymbol{\theta}}^{(t)}). \quad (42)$$

Moreover, the computation of the update is simple. To obtain $\mathbf{e}^{(t)}$ we apply the inverse UDWT to $\boldsymbol{\theta} = (\mathbf{0}, \hat{\boldsymbol{\theta}}_I^{(t+1)})$ to obtain $\mathbf{W}_I \hat{\boldsymbol{\theta}}_I^{(t)}$. This can be computed in $O(N \log N)$ operations. Finally, notice that (38) is simply a standard DWT denoising operation (with the threshold/shrinkage function associated with $pen(\cdot)$) applied to $\mathbf{e}^{(t)}$, which can be computed in $O(N)$ operations.

Being a GEM algorithm, this method has all the monotonicity guarantees of EM and is thus of theoretical interest. However, it turns out that, in all the experiments carried out, this approach performs slightly worse than the two previous methods; for this reason, we will not further consider it in this paper.

IX. EXPERIMENTAL RESULTS

In this section, we present a set of experimental results illustrating the performance of the proposed approach and comparing it with some state-of-the-art methods recently described in [17], [22], and [29]. We consider only the TI versions of the algorithm: the UDWT-based method (using the UDWT filterbank of [21]) and the method based on random shifts; the reason for this choice is that the TI versions clearly and consistently outperform those that use the orthogonal DWT. Moreover, we do not consider the noise-adaptive version described in Section VI; this is because we always achieve better performance using a fixed noise variance, which can be easily estimated directly from the observed image using the MAD scheme proposed in [10].

In all the experiments, we employ Daubechies-2 (Haar) wavelets; we have tried other wavelets and the results are always very similar to those obtained with Daubechies-2. The algorithm is initialized with a Wiener estimate, as given by (5), with $\boldsymbol{\mu} = \mathbf{0}$ and $\mathbf{G} = 10^3 \mathbf{I}$, and the convergence criterion is

$$\frac{\|\hat{\mathbf{x}}^{(t+1)} - \hat{\mathbf{x}}^{(t)}\|_2}{\|\hat{\mathbf{x}}^{(t)}\|_2} < \delta,$$

where δ is a threshold, typically set to $10^{-3} \sigma^2$. As discussed in Section V-E, we set $\alpha = \sigma$; we found experimentally that this is a good general-purpose choice.

In the first set of tests, we consider the setup of [29] and [3]: uniform blur of size 9×9 , and the noise variance such that the SNR of the noisy image, with respect to the blurred image without

TABLE I

SNR IMPROVEMENTS (SNRI) OBTAINED BY SEVERAL VARIANTS OF THE PROPOSED ALGORITHM ON THE BLURRED IMAGE SHOWN IN FIGURE 2

Method	SNRI
Rule (24), UDWT	7.47dB
Rule (24), random shifts	7.59dB
Modified Laplacian, UDWT	7.26dB
Modified Laplacian, random shifts	7.34dB
Soft-threshold, UDWT	7.26dB
Soft-threshold, random shifts	6.33dB
Result by Neelamani <i>et al</i> [29]	7.3dB
Result by Banham and Katsaggelos [3]	6.7dB

noise (BSNR), is $40dB$ (this corresponds to $\sigma^2 \simeq 0.308$). We have restored this image using six variants of the algorithm: the denoising rule (24), the rule corresponding to the modified Laplacian prior (with $\eta = 0.35$ and $\beta = 0.02$, see (29)), and the soft-threshold rule, each with the UDWT-based method and the random shifts scheme. The SNR improvements obtained by the several algorithms are summarized in Table I, showing that our methods perform competitively (some versions better, others slightly worse) than the one in [29]. Figure 2 shows the original, blurred/noisy, and restored images, using rule (24) and the UDWT-based method. The other restored images are visually indistinguishable from this one, so we do not show them here. Finally, in Figure 3, we plot the evolution of the SNR improvement along the EM algorithm, for the UDWT-based and the random shifts algorithm, both with rule (24) (the other versions of the algorithm evolve similarly). We can observe that convergence is obtained after $200 \sim 300$ iterations.

In the second set of tests, we replicate the experimental condition of [17]. The point spread function of the blur operator is given by $h_{ij} = (1 + i^2 + j^2)^{-1}$, for $i, j = -7, \dots, 7$. Noise



Fig. 2. Original image (top), blurred image (bottom left) and restored image (bottom right) using the UDWT version of our algorithm with rule (24).

variances considered are $\sigma^2 = 2$ and $\sigma^2 = 8$. The SNR improvements obtained are summarized in Table II, together with the results reported in [17]. Figure 4 shows the original image, the two blurred/noisy images, and the corresponding restorations, obtained with rule (24) and the UDWT-based method. The SNR improvements obtained by our method are very similar to those reported in [17]; notice that [17] uses a more sophisticated wavelet transform and prior model. Finally, in Figure 5, we plot the evolution of the SNR improvement along the EM iterations, for the UDWT-based algorithm with rule (24) (the other versions of the algorithm evolve similarly). we see that convergence is achieved after approximately 40 and $8 \sim 10$ iterations, respectively, for $\sigma^2 = 2$ and $\sigma^2 = 8$.




Fig. 3. SNR improvement along the iterations of EM (left: UDWT-based method; right: random shifts), for the example of Figure 2.

In the final set of tests we have used the blur filter and noise variance considered in [22]. Specifically, the original image was blurred by a 5×5 separable filter with weights [14641]/16 (in both horizontal and vertical directions) and then contaminated with white Gaussian noise of standard deviation $\sigma = 7$. The SNR improvements obtained by the six instances of our algorithm are reported in Table III. The original, blurred, and restored images are shown in Figure 6. In this case, convergence is obtained after $5 \sim 7$ iterations.

TABLE II

SNR IMPROVEMENTS OBTAINED BY SEVERAL VARIANTS OF THE PROPOSED ALGORITHM ON THE IMAGES SHOWN IN FIGURE 4

Method	$\sigma^2 = 2$	$\sigma^2 = 8$
Rule (24), UDWT	6.91dB	4.88dB
Rule (24), random shifts	6.93dB	4.37dB
Modified Laplacian, UDWT	6.39dB	4.51dB
Modified Laplacian, random shifts	6.33dB	4.22dB
Soft-threshold, UDWT	6.36dB	4.12dB
Soft-threshold, random shifts	6.42dB	4.01dB
Results by Jalobeanu <i>et al</i> [17]	6.75dB	4.85dB



Fig. 4. Blurred and noisy images (top row), with $\sigma^2 = 2$ (left) and $\sigma^2 = 8$ (right), and corresponding restored images (bottom row).

We can observe a clear trend in the behavior of the algorithm: for larger noise variance, convergence is achieved in fewer iterations (recall from the results above: 200 \sim 300 iterations for $\sigma^2 \simeq 0.308$; \sim 40 iterations for $\sigma^2 = 2$; 8 \sim 10 iterations for $\sigma^2 = 8$; and 5 \sim 7 iterations for $\sigma^2 = 7^2$). As the number of iterations decreases, the performance of the random-shifts-based method degrades, since it does not cover enough shifts to achieve approximate shift-invariance.

X. CONCLUSIONS

This paper proposed a wavelet-based MPLE/MAP criterion for image deconvolution. The estimate must be computed numerically, and we derived an EM algorithm for this purpose, leading to a simple procedure that alternate between Fourier domain filtering and wavelet domain




Fig. 5. SNR improvement along the iterations of EM for the example of Figure 4 (left, $\sigma^2 = 2$; right, $\sigma^2 = 8$).

TABLE III

SNR IMPROVEMENTS OBTAINED BY SEVERAL VARIANTS OF THE PROPOSED ALGORITHM ON THE
BLURRED IMAGE SHOWN IN FIGURE 6

Method	SNRI
Rule (24), UDWT	2.94dB
Rule (24), random shifts	1.71dB
Modified Laplacian, UDWT	2.75dB
Modified Laplacian, random shifts	1.77dB
Soft-threshold, UDWT	2.75dB
Soft-threshold, random shifts	1.61dB
Best result by Liu and Moulin [22]	1.078dB

denoising. We have also proposed extensions of the algorithm which perform shift-invariance restoration. Experimentally, our approach performs competitively with two of the best existing methods. Our analysis sheds light on the nature of wavelet-based image restoration; in particular this is the first work we are aware of that carefully investigates and describes the subtle distinctions between the application of orthogonal and (non-orthogonal) translation-invariant DWTs to image restoration.



Fig. 6. Original image (left), blurred image (middle) and restored image (right) using the noise-adaptive version of our algorithm with rule (24).

ACKNOWLEDGMENTS

The authors would like to thank Jeff Fessler and André Jalobeanu for their helpful discussions and insightful comments that helped to improve this paper.

REFERENCES

- [1] H. Andrews and B. Hunt. *Digital Image Restoration*, Prentice Hall, Englewood Cliffs, NJ, 1977.
- [2] A. Antoniadis and J. Fan. “Regularized wavelet approximations,” *Journal of the American Statistical Association*, vol. 96, pp. 939-967, 2001.
- [3] M. Banham and A. Katsaggelos. “Spatially adaptive wavelet-based multiscale image restoration,” *IEEE Trans. on Image Processing*, vol. 5, pp. 619-634, 1996.
- [4] M. Belge, M. E. Kilmer, and E. L. Miller, “Wavelet domain image restoration with adaptive edge-preserving regularity,” *IEEE Trans. Image Processing*, vol. 9, no. 4, pp. 597–608, 2000.
- [5] R. Coifman and D. Donoho. “Translation invariant de-noising,” in A. Antoniadis and G. Oppenheim, editors, *Wavelets and Statistics*, Lecture Notes in Statistics vol. 103, pp. 125–150, Springer-Verlag, New York, 1995.
- [6] M. Crouse, R. Nowak, and R. Baraniuk. “Wavelet-based statistical signal processing using hidden Markov models,” *IEEE Trans. on Signal Processing*, vol. 46, pp. 886–902, 1998.
- [7] A. Dempster, N. Laird, and D. Rubin. “Maximum likelihood estimation from incomplete data via the EM algorithm,” *Journal of the Royal Statistical Society B*, vol. 39, pp. 1–38, 1977.
- [8] P. de Rivaz and N. Kingsbury. “Bayesian image deconvolution and denoising using complex wavelets,” in *Proc. of the IEEE Intern. Conf. on Image Proc. – ICIP’2001*, Thessaloniki, Greece, 2001.

- [9] D. Donoho. “Nonlinear solution of linear inverse problems by wavelet-vaguelette decompositions,” *Journal of Applied and Computational Harmonic Analysis*, vol. 1, pp. 100–115, 1995.
- [10] D. L. Donoho and I. M. Johnstone, “Adapting to unknown smoothness via wavelet shrinkage,” *Journal of the American Statistical Association*, vol. 90, no. 432, pp. 1200–1224, 1995.
- [11] J. Fessler, A. Hero. “Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms,” *IEEE Trans. on Image Processing*, vol. 4, pp. 1417–29, 1995.
- [12] M. Figueiredo and J. Leitão, “Unsupervised Image Restoration and Edge Location using Compound Gauss-Markov random fields and the MDL principle”, *IEEE Trans. on Image Processing*, vol. 6, pp. 1089–1102, 1997.
- [13] M. Figueiredo and R. Nowak. “Bayesian wavelet-based signal estimation using non-informative priors,” in *Proc. of the 32nd Asilomar Conference on Signals, Systems, and Computers*, pp. 1368 – 1373, Monterey, CA, 1998.
- [14] M. Figueiredo and R. Nowak. “Wavelet-based image estimation: an empirical Bayes approach using Jeffreys’ noninformative prior,” *IEEE Trans. on Image Processing*, vol. 10, pp. 1322–1331, 2001.
- [15] H. Gazzah and P. Regalia and J.-P. Delmas. “Asymptotic eigenvalue distribution of block Toeplitz matrices and application to blind SIMO channel identification,” *IEEE Trans. on Information Theory*, vol. 47, pp. 1243–1251, 2001.
- [16] A. Jain. *Fundamentals of Digital Image Processing*, Prentice Hall, Englewood Cliffs, N. J., 1989.
- [17] A. Jalobeanu, N. Kingsbury, and J. Zerubia. “Image deconvolution using hidden Markov tree modeling of complex wavelet packets,” in *Proc. of the IEEE Intern. Conf. on Image Proc. – ICIP’2001*, Thessaloniki, Greece, 2001.
- [18] J. Kalifa, S. Mallat, and B. Rougé. “Deconvolution by thresholding in mirror wavelet bases,” in *IEEE Trans. on Image Processing*, vol. 12, 2003 (to appear).
- [19] R. Lagendijk, J. Biemond, D. Boekee. “Identification and restoration of noisy blurred images using the expectation-maximization algorithm,” *IEEE Trans. on Acoustics, Speech, and Signal Processing*, vol. 38, pp. 1180–1191, 1990.
- [20] L. Landweber. “An iterative formula for Fredholm integral equations of the first kind,” *American Journal of Mathematics*, vol. 73, pp. 615–624, 1951.
- [21] M. Lang, H. Guo, J. E. Odegard, C. S. Burrus and R. O. Wells. “Noise reduction using an undecimated discrete wavelet transform,” *IEEE Signal Processing Letters*, vol. 3, pp. 10–12, 1996.
- [22] J. Liu and P. Moulin. “Complexity-Regularized Image Restoration,” *Proc. IEEE Int. Conf. on Image Proc. (ICIP’98)*, Vol. 1, pp. 555–559, Chicago, Oct. 1998.
- [23] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego, 1998.

- [24] G. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. John Wiley & Sons, New York, 1997.
- [25] M. Mihçak, I. Kozintsev, K. Ramchandran, and P. Moulin. “Low-complexity image denoising based on statistical modeling of wavelet coefficients,” *Signal Processing Letters*, vol. 6, pp. 300–303, 1999.
- [26] P. Moulin and J. Liu. “Analysis of multiresolution image denoising schemes using generalized - Gaussian and complexity priors,” *IEEE Trans. on Information Theory*, vol. 45, pp. 909–919, 1999.
- [27] P. Moulin and J. Liu. “Statistical imaging and complexity regularization,” *IEEE Trans. on Information Theory*, vol. 46, pp. 1881–1895, 2000.
- [28] F. Murtagh, J. Starck and A. Bijaoui, *Image Processing and Data Analysis*. Cambridge Univeristy Press, 1998.
- [29] R. Neelamani, H. Choi, and R. Baraniuk. “Wavelet-based deconvolution for ill-conditioned systems,” *IEEE Trans. on Image Processing*, 2001 (submitted).
- [30] R. Nowak and D. Kolaczyk. “A statistical multiscale framework for Poisson inverse problems,” *IEEE Trans. on Information Theory*, vol. 46, pp. 1811–1825, 2000.
- [31] L. Scharf. *Statistical Signal Processing*. Addison Wesley Publishing Company, Reading, Massachusetts, 1991.
- [32] O. Strand. “Theory and methods related to the singular-function expansion and Landweber’s iteration for integral equations of the first kind,” *SIAM Journal of Numerical Analysis*, vol. 11, pp. 798–825, 1974.
- [33] Y. Wan and R. Nowak. “A wavelet-based approach to joint image restoration and edge detection,” in *SPIE Conference on Wavelet Applications in Signal and Image Processing VII*, Denver, CO, 1999. SPIE Vol. 3813.
- [34] J. Z. G. Wang and G.-W. Pan, “Solution of inverse problem in image processing by wavelet expansion,” *IEEE Trans. Image Processing*, vol. 4, no. 5, pp. 579–593, 1995.
- [35] C. Wu. “On the convergence properties of the EM algorithm,” *Annals of Statistics*, vol. 11, pp. 95-103, 1983.