

AN EMPIRICAL BAYES EM-WAVELET UNIFICATION FOR SIMULTANEOUS DENOISING, INTERPOLATION, AND/OR DEMOSAICING

Keigo Hirakawa and Xiao-Li Meng

Harvard University
Department of Statistics
1 Oxford Street, Cambridge, MA 02138
{hirakawa,meng}@stat.harvard.edu

ABSTRACT

We present a unified framework for coupling the EM algorithm with the Bayesian hierarchical modeling of neighboring wavelet coefficients of image signals. Within this framework, problems with missing pixels or pixel components, and hence unobservable wavelet coefficients, are handled simultaneously with denoising. The hyper-parameters of the model are estimated via the marginal likelihood by the EM algorithm, and a part of the output of its E-step automatically provide optimal estimates, given the specified Bayesian model, of the noise-free image. This unified empirical-Bayes based framework, therefore, offers a statistically principled and extremely flexible approach to a wide range of pixel estimation problems including image denoising, image interpolation, demosaicing, or any combinations of them.

Index Terms— wavelets, missing data, denoising, interpolation

1. INTRODUCTION

The statistical modeling of image signals in the wavelet domain is primarily based on the observations that a wavelet coefficient and its adjacent/parent coefficients of image signals are correlated. Here we use Bayesian hierarchical modeling to capture such correlations. While similar approaches have been suggested [1, 2, 3], wavelet-based imaging algorithms commonly use heuristically determined parameters, which are often restrictive or inaccurate. To overcome these shortcomings, we assume a general model form and couple the EM algorithm framework with the Bayesian models to estimate the hyper-parameters via the marginal likelihood, that is, we adopt the empirical Bayes approach. One key benefit of this approach is that a part of the output provides automatically provides the optimal Bayesian estimate (given the specified model) of the noise-free image, as specified. Under the assumption of additive noise, a similar EM algorithm was designed in [4, 5], but the parameter estimation is limited to that of noise correlation matrix.

Another key benefit of this approach is that it can easily be extended to *simultaneously* deal with problems with missing or incomplete pixel values, either because of mechanical designs (e.g., demosaicing) or because distortion (e.g., picture impainting) or both. In the context of wavelet-based image processing, missing or incomplete pixel is a difficult problem because wavelet transform takes a linear combination of image signal, and thus many, or even *all* of the noisy wavelet coefficients are unobserved. Using the EM algorithm, however, the considerations for the missing pixel imputation and the model parameter estimation can be combined into a unified theoretical framework. The E-step imputation can be viewed as a solution

Authors thank Dr. J. Portilla and Dr. E.P. Simoncelli for their helpful comments, and NSF for partial funding.

to a model-based image denoising problem with incomplete data, or a model-based image interpolation given noisy data. This generalization has wide ranges of useful applications that include image denoising, image interpolation, demosaicing, or *combinations of them*.

Given the space limitation and the considerable generality of our approach, we assume the familiarity with the EM algorithm [6] and focus on the basic theory under additive Gaussian noise. However, neither the additivity nor the normality is essential to our general approach—they just simplify both the presentation and implementation. We present the Bayesian hierarchical model in Section 2, and its EM algorithm in Section 3. We illustrate the method with experimental results in Section 4, and conclude in Section 5.

2. BAYESIAN HIERARCHICAL MODELING

Suppose there are N pixels in an image $\mathbf{f}^\top = [f_1, \dots, f_N]$, where f_i is the true pixel value at location $i \in \{1, \dots, N\}$. In an image denoising problem, the observed image, \mathbf{y} is corrupted by noise. Assuming an independent additive Gaussian noise, let $\mathbf{y} = \mathbf{f} + \mathbf{e}$, where $\mathbf{e} \sim \mathcal{N}(0, \Sigma_e)$. The objective here is to estimate \mathbf{f} given \mathbf{y} . Suppose only some of the pixels in the image \mathbf{y} are observed. Define $\mathbf{y}_{\text{obs}}^\top = [y_{\text{obs}1}, \dots, y_{\text{obs}N_o}]$ and $\mathbf{y}_{\text{mis}}^\top = [y_{\text{mis}1}, \dots, y_{\text{mis}N_m}]$ as observed and missing pixel values, respectively, where $N_o + N_m = N$ and, without loss of generality, $\mathbf{y}^\top = [\mathbf{y}_{\text{obs}}^\top, \mathbf{y}_{\text{mis}}^\top]$. We represent \mathbf{f} and \mathbf{y} by wavelet transform, $\mathbf{d} = \mathbf{W}\mathbf{f}$ and $\mathbf{w} = \mathbf{W}\mathbf{y}$, respectively. Here, \mathbf{W} is the two-dimensional wavelet transform matrix operator with M subbands. Define w_{km} as the wavelet coefficient at the k -th pixel location in the m -th subband ($m = 0$ represents the scaling coefficients). Define vector $\mathbf{w}_{km} = \{w_{k'm'}\} \in \mathbb{R}^n$, where $w_{k'm'}$ are coefficients in \mathbf{w} that are in the *neighborhood* of w_{km} (to be made precise in the sequel), and let \mathbf{d}_{km} be defined similarly. Let \mathbf{W}_{km} correspond to the appropriate rows of \mathbf{W} such that $\mathbf{w}_{km} = \mathbf{W}_{km}\mathbf{y}$. Here we assume orthogonal wavelet transform (i.e. $\mathbf{W}\mathbf{W}^\top = \mathbf{I}$), but the theoretical framework extends to non-orthogonal and over-complete transforms in a straightforward manner (see Section 4).

The Bayesian hierarchical model of the image wavelet coefficients considered in this paper is based on ideas exploited in [1]. Assuming signal-independent additive Gaussian noise, wavelet coefficients are modeled hierarchically. Motivated by the observation that wavelet coefficients of natural images tend toward heavy-tailed distributions (see [1]), we model \mathbf{d}_{km} by a t -distribution with ν_m degrees of freedom. That is, \mathbf{w}_{km} and \mathbf{d}_{km} take the following form:

$$\mathbf{w}_{km} | \mathbf{d}_{km} \sim \text{indep } \mathcal{N}(\mathbf{d}_{km}, \Sigma_{wm}) \quad (1)$$

$$\mathbf{d}_{km} | q_{km} \sim \text{indep } \mathcal{N}(\boldsymbol{\xi}_m, \Sigma_{dm} / q_{km}) \quad (2)$$

$$q_{km} \sim \text{indep } \chi_{\nu_m}^2 / \nu_m. \quad (3)$$

The use of the unobserved mean chi-squared variables $\mathbf{q} = \{q_{km}\}$ to represent the t_{w_m} as part of the *augmented data* is a well-known strategy for easy EM implementation [7]. Here our goal is to estimate the (hyper-)parameter $\theta = \{\Sigma_{wm}, \Sigma_{dm}, \nu_m, \xi_m\}$ by maximizing the marginal log-likelihood $\ell(\theta|\mathbf{y}_{\text{obs}}) = \log p(\mathbf{y}_{\text{obs}}|\theta)$. The direct maximization is very difficult because of the missing pixel values. The EM algorithm circumvents this problem by iteratively maximizing the much easier augmented-data log-likelihood $\ell(\theta|\mathbf{x}) = \log p(\mathbf{x}|\theta)$, where $\mathbf{x} = \{\mathbf{y}, \mathbf{f}, \mathbf{q}\}$, or equivalently $\{\mathbf{w}, \mathbf{d}, \mathbf{q}\}$ are the *augmented data*. Because only \mathbf{y}_{obs} is observed, the E-step computes the imputation of augmented-data sufficient statistics, needed by the M-step, which includes $\hat{\mathbf{d}}^{[t+1]} = E[\mathbf{d}|\mathbf{y}_{\text{obs}}, \theta^{[t]}]$. This can be used to compute the ‘‘best’’ estimate, at the $(t+1)$ st iteration, of \mathbf{f} , the noise-free image, as $\hat{\mathbf{f}}^{[t+1]} = \mathbf{W}^T \hat{\mathbf{d}}^{[t+1]}$.

Before we proceed, we point out a potential theoretical incompatibility inherited in (1)-(2) when the neighborhoods are allowed to have overlaps. As we have no space to discuss this issue in any details, in particularly its connection with the over-complete expansions, we simply point out that assuming independence between overlapping neighborhoods provides a computationally efficient approximation for practical implementation.

3. EM ALGORITHM DERIVATION

Given $\theta^{[t]}$, the $(t+1)$ st iteration of the EM algorithm first calls for

$$\begin{aligned} Q(\theta; \theta^{[t]}) &= E \left[\log p(\mathbf{x}|\theta) \middle| \mathbf{y}_{\text{obs}}, \theta^{[t]} \right] \\ &= \sum_{k,m} E \left[\log p(\mathbf{w}_{km}|\mathbf{d}_{km}; \Sigma_{wm}) + \log p(\mathbf{d}_{km}|\mathbf{q}_{km}; \xi_m, \Sigma_{dm}) \right. \\ &\quad \left. + \log p(q_{km}|\nu_m) \middle| \mathbf{y}_{\text{obs}}, \theta^{[t]} \right]. \end{aligned}$$

A celebrated result of EM algorithm [6] states that the choice of θ that maximizes $Q(\theta; \theta^{[t]})$, that is, the next iterate $\theta^{[t+1]}$ increases our objective function: $\ell(\theta^{[t+1]}|\mathbf{y}_{\text{obs}}) \geq \ell(\theta^{[t]}|\mathbf{y}_{\text{obs}})$. The advantage of maximizing $Q(\theta; \theta^{[t]})$ instead of directly the targeted $\ell(\theta|\mathbf{y}_{\text{obs}})$ is that it is often the case that $\log p(\mathbf{x}|\theta)$ depends on the missing part of \mathbf{x} linearly via a set of *augmented-data sufficient statistics* $\mathbf{S}(\mathbf{x})$. Consequently, maximizing $Q(\theta; \theta^{[t]})$ is the same as maximizing $\ell(\theta|\mathbf{x})$, but with the (partially) missing $\mathbf{S}(\mathbf{x})$ imputed via the E-step: $\mathbf{S}^{[t+1]}(\mathbf{x}) = E[\mathbf{S}(\mathbf{x})|\mathbf{y}_{\text{obs}}, \theta^{[t]}]$.

For our problem, this is most clearly seen when ν_m is known, and thus the $\log p(q_{km}|\nu_m)$ part is unneeded for $Q(\theta; \theta^{[t]})$. We therefore present our basic algorithm first under this assumption. It is then easy to verify that $\mathbf{S}(\mathbf{x}) = \{\mathbf{S}_1(\mathbf{w}, \mathbf{d}), \mathbf{S}_2(\mathbf{q}, \mathbf{d})\}$, where

$$\begin{aligned} \mathbf{S}_1(\mathbf{w}, \mathbf{d}) &= \{\mathbf{w}_{km} \mathbf{w}_{km}^\top, \mathbf{d}_{km} \mathbf{d}_{km}^\top, \mathbf{w}_{km} \mathbf{d}_{km}^\top\}, \\ \mathbf{S}_2(\mathbf{q}, \mathbf{d}) &= \{q_{km}, q_{km} \mathbf{d}_{km}, q_{km} \mathbf{d}_{km} \mathbf{d}_{km}^\top\}. \end{aligned}$$

Indeed, since θ is reduced to $\theta_1 = \{\Sigma_{wm}, \Sigma_{dm}, \xi_m\}$, the maximizer of $Q(\theta_1; \theta_1^{[t]})$ is the weighted least-square estimate [6]:

$$\xi_m^{[t+1]} = \frac{\sum_k E[q_{km} \mathbf{d}_{km} | \mathbf{y}_{\text{obs}}, \theta^{[t]}]}{\sum_k E[q_{km} | \mathbf{y}_{\text{obs}}, \theta^{[t]}]}, \quad (4)$$

$$\Sigma_{wm}^{[t+1]} = n_m^{-1} \sum_k E[\mathbf{w}_{km} \mathbf{w}_{km}^\top - \mathbf{w}_{km} \mathbf{d}_{km}^\top - \mathbf{d}_{km} \mathbf{w}_{km}^\top + \mathbf{d}_{km} \mathbf{d}_{km}^\top | \mathbf{y}_{\text{obs}}, \theta^{[t]}], \quad (5)$$

$$\Sigma_{dm}^{[t+1]} = n_m^{-1} \sum_k E[q_{km} (\mathbf{d}_{km} \mathbf{d}_{km}^\top - \mathbf{d}_{km} \xi_m^{[t+1]\top} - \xi_m^{[t+1]} \mathbf{d}_{km}^\top + \xi_m^{[t+1]} \xi_m^{[t+1]\top}) | \mathbf{y}_{\text{obs}}, \theta^{[t]}], \quad (6)$$

where n_m is the number of wavelet coefficients in the m -th subband.

3.1. A Conditional E-step

To carry out the above calculation, we need to compute $\mathbf{S}^{[t+1]}(\mathbf{x}) = E[\mathbf{S}(\mathbf{x})|\mathbf{y}_{\text{obs}}, \theta^{[t]}]$. This computation involves the t -distribution for \mathbf{q} , and therefore is complicated. However, when conditioned on \mathbf{q} , only multivariate normal distributions are involved. So our general strategy is to first take the expectation conditioned on \mathbf{q} (and \mathbf{y}_{obs} with $\theta = \theta^{[t]}$), and then integrate out \mathbf{q} numerically when we complete the M-step.

Specifically, let $\varphi = E(\mathbf{y}|\mathbf{q}; \theta) = \mathbf{W}^T E(\mathbf{w}|\mathbf{q})$ and $\Omega_{|\mathbf{q}} = \text{Cov}(\mathbf{y}|\mathbf{q}, \theta) = \mathbf{W}^\top \text{Cov}(\mathbf{w}|\mathbf{q}, \theta) \mathbf{W}$ be the conditional mean and variance of \mathbf{y} given \mathbf{q} under our model (1)-(3). Note that φ is free of \mathbf{q} under our model, and given \mathbf{q} , $\Omega_{|\mathbf{q}}$ and φ can be defined completely by the parameters in θ . Define φ_{obs} the sub-vector of φ that corresponds to \mathbf{y}_{obs} , and partition $\text{Cov}(\mathbf{y}|\mathbf{q}, \theta)$ into four submatrices according to the partition $\mathbf{y} = (\mathbf{y}_{\text{obs}}, \mathbf{y}_{\text{mis}})$ as $\Omega_{|\mathbf{q}} = [\Omega_{o|q}, \Omega_{om|q}; \Omega_{mo|q}, \Omega_{m|q}]$. Conditioned on \mathbf{q} , the noisy image \mathbf{y} is imputed via the standard regression estimate:

$$\hat{\mathbf{y}}_{|\mathbf{q}} = E(\mathbf{y}|\mathbf{q}, \mathbf{y}_{\text{obs}}, \theta) = \varphi + [\Omega_{o|q}; \Omega_{mo|q}] \Omega_{o|q}^{-1} (\mathbf{y}_{\text{obs}} - \varphi_{\text{obs}}).$$

And $\text{Cov}(\mathbf{y}_{\text{mis}}|\mathbf{q}, \mathbf{y}_{\text{obs}}, \theta)$ is a Schur complement of $\Omega_{o|q}$ in $\Omega_{|\mathbf{q}}$:

$$\text{Cov}(\mathbf{y}_{\text{mis}}|\mathbf{q}, \mathbf{y}_{\text{obs}}, \theta) = \Omega_{m|q} - \Omega_{mo|q} \Omega_{o|q}^{-1} \Omega_{om|q}.$$

Imputation of other data vectors follow easily:

$$\begin{aligned} \hat{\mathbf{w}}_{km|q} &= E[\mathbf{w}_{km}|\mathbf{q}, \mathbf{y}_{\text{obs}}, \theta] = \mathbf{W}_{km} E(\mathbf{y}|\mathbf{q}, \mathbf{y}_{\text{obs}}, \theta) = \mathbf{W}_{km} \hat{\mathbf{y}}_{|\mathbf{q}} \\ \hat{\mathbf{d}}_{km|q} &= E[\mathbf{d}_{km}|\mathbf{q}, \mathbf{y}_{\text{obs}}, \theta] \\ &= E[E[\mathbf{d}_{km}|\mathbf{w}_{km}, \mathbf{q}, \mathbf{y}_{\text{obs}}, \theta] | \mathbf{q}, \mathbf{y}_{\text{obs}}, \theta] \\ &= E[\xi_m + (\Sigma_{wm}^{-1} + q_{km} \Sigma_{dm}^{-1})^{-1} \Sigma_{wm}^{-1} (\mathbf{w}_{km} - \xi_m) | \mathbf{q}, \mathbf{y}_{\text{obs}}, \theta] \\ &= \xi_m + (\Sigma_{wm}^{-1} + q_{km} \Sigma_{dm}^{-1})^{-1} \Sigma_{wm}^{-1} (\hat{\mathbf{w}}_{km|q} - \xi_m). \end{aligned}$$

Using the formulae above, all the covariance matrices (conditioned on \mathbf{q}) that are needed to carry out $E[\mathbf{S}(\mathbf{x})|\mathbf{q}, \mathbf{y}_{\text{obs}}, \theta]$ can be calculated analytically. See (7), (8), and (9), where \mathbf{I} is an identity matrix, and $\mathbf{C}_{dkm|q} = (\Sigma_{wm}^{-1} + q_{km} \Sigma_{dm}^{-1})^{-1}$.

3.2. The M-step

By (4)-(6), the M-step updates $\theta_1 = \{\Sigma_{wm}, \Sigma_{dm}, \xi_m\}$ by

$$\xi_m^{[t+1]} = \frac{\sum_k \int_0^\infty q_{km} \hat{\mathbf{d}}_{km|q} p(\mathbf{q}|\mathbf{y}_{\text{obs}}, \theta^{[t]}) d\mathbf{q}}{\sum_k \int_0^\infty q_{km} p(\mathbf{q}|\mathbf{y}_{\text{obs}}, \theta^{[t]}) d\mathbf{q}},$$

$$\Sigma_{wm}^{[t+1]} = n_m^{-1} \sum_k \int_0^\infty (\Lambda_{wkm|q} - \Lambda_{wdkm|q} - \Lambda_{wdkm|q}^\top + \Lambda_{dkm|q}) p(\mathbf{q}|\mathbf{y}_{\text{obs}}, \theta^{[t]}) d\mathbf{q},$$

$$\Sigma_{dm}^{[t+1]} = n_m^{-1} \sum_k \int_0^\infty q_{km} (\Lambda_{dkm|q} - \hat{\mathbf{d}}_{km|q} \xi_m^{[t+1]\top} - \xi_m^{[t+1]} \hat{\mathbf{d}}_{km|q}^\top + \xi_m^{[t+1]} \xi_m^{[t+1]\top}) p(\mathbf{q}|\mathbf{y}_{\text{obs}}, \theta^{[t]}) d\mathbf{q}.$$

Likewise, the estimate for the noise-free wavelet coefficients are found by integration:

$$\hat{\mathbf{d}}_{km} = E[\mathbf{d}_{km}|\mathbf{y}_{\text{obs}}, \theta] = \int_0^\infty \hat{\mathbf{d}}_{km|q} p(\mathbf{q}|\mathbf{y}_{\text{obs}}, \theta) d\mathbf{q}.$$

The integrals above can be simplified via the Bayes rule. Specifically, for any integrable $g(\mathbf{q})$, we have

$$E[g(\mathbf{q})|\mathbf{y}_{\text{obs}}, \theta] = \frac{\int_0^\infty g(\mathbf{q}) \phi_{\text{obs}}(\mathbf{q}) \prod_{k,m} \tilde{p}(q_{km}|\nu_m) d\mathbf{q}_{km}}{\int_0^\infty \phi_{\text{obs}}(\mathbf{q}) \prod_{k,m} \tilde{p}(q_{km}|\nu_m) d\mathbf{q}_{km}}$$

$$\begin{aligned}\Lambda_{wkm|q} &= E[\mathbf{w}_{km} \mathbf{w}_{km}^\top | \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}] = \mathbf{W}_{km} E[\mathbf{y} \mathbf{y}^\top | \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}] \mathbf{W}_{km}^\top \\ &= \mathbf{W}_{km} \left[\text{Cov}(\mathbf{y} | \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}) + \hat{\mathbf{y}} | \mathbf{q} \hat{\mathbf{y}}^\top | \mathbf{q} \right] \mathbf{W}_{km}^\top = \mathbf{W}_{km} \text{Cov}(\mathbf{y} | \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}) \mathbf{W}_{km}^\top + \hat{\mathbf{w}}_{km|q} \hat{\mathbf{w}}_{km|q}^\top\end{aligned}\quad (7)$$

$$\begin{aligned}\Lambda_{dkm|q} &= E[\mathbf{d}_{km} \mathbf{d}_{km}^\top | \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}] = E \left\{ E[\mathbf{d}_{km} \mathbf{d}_{km}^\top | \mathbf{w}, \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}] \middle| \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta} \right\} \\ &= E \left\{ \text{Cov}(\mathbf{d}_{km} | \mathbf{w}, \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}) + E[\mathbf{d}_{km} | \mathbf{w}, \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}] E[\mathbf{d}_{km}^\top | \mathbf{w}, \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}] \middle| \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta} \right\} \\ &= \mathbf{C}_{dkm|q} + E \left\{ [\boldsymbol{\xi}_m + \mathbf{C}_{dkm|q} \boldsymbol{\Sigma}_{wm}^{-1} (\mathbf{w}_{km} - \boldsymbol{\xi}_m)] \left[(\mathbf{w}_{km} - \boldsymbol{\xi}_m)^\top \boldsymbol{\Sigma}_{wm}^{-1} \mathbf{C}_{dkm|q} + \boldsymbol{\xi}_m^\top \right] \middle| \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta} \right\} \\ &= \mathbf{C}_{dkm|q} + (\mathbf{I} - \mathbf{C}_{dkm|q} \boldsymbol{\Sigma}_{wm}^{-1}) (\boldsymbol{\xi}_m \boldsymbol{\xi}_m^\top) (\mathbf{I} - \boldsymbol{\Sigma}_{wm}^{-1} \mathbf{C}_{dkm|q}) + (\mathbf{I} - \mathbf{C}_{dkm|q} \boldsymbol{\Sigma}_{wm}^{-1}) (\boldsymbol{\xi}_m \hat{\mathbf{w}}_{km|q}^\top) (\boldsymbol{\Sigma}_{wm}^{-1} \mathbf{C}_{dkm|q}) \\ &\quad + (\mathbf{C}_{dkm|q} \boldsymbol{\Sigma}_{wm}^{-1}) (\hat{\mathbf{w}}_{km|q} \boldsymbol{\xi}_m^\top) (\mathbf{I} - \boldsymbol{\Sigma}_{wm}^{-1} \mathbf{C}_{dkm|q}) + (\mathbf{C}_{dkm|q} \boldsymbol{\Sigma}_{wm}^{-1}) \Lambda_{wkm|q} (\boldsymbol{\Sigma}_{wm}^{-1} \mathbf{C}_{dkm|q})\end{aligned}\quad (8)$$

$$\begin{aligned}\Lambda_{wdkm|q} &= E[\mathbf{w}_{km} \mathbf{d}_{km}^\top | \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}] = E \left\{ \mathbf{w}_{km} E[\mathbf{d}_{km}^\top | \mathbf{w}, \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}] \middle| \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta} \right\} \\ &= E \left\{ \mathbf{w}_{km} \left[(\mathbf{w}_{km} - \boldsymbol{\xi}_m)^\top \boldsymbol{\Sigma}_{wm}^{-1} \mathbf{C}_{dkm|q} + \boldsymbol{\xi}_m^\top \right] \middle| \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta} \right\} = \Lambda_{wkm|q} \boldsymbol{\Sigma}_{wm}^{-1} \mathbf{C}_{dkm|q} + \hat{\mathbf{w}}_{km|q} \boldsymbol{\xi}_m^\top (\mathbf{I} - \boldsymbol{\Sigma}_{wm}^{-1} \mathbf{C}_{dkm|q})\end{aligned}\quad (9)$$

where, denoting $|\cdot|$ the determinant operator,

$$\begin{aligned}\phi_{\text{obs}}(\mathbf{q}) &= |\Omega_{o|q}|^{-\frac{1}{2}} \exp \left\{ -\frac{(\mathbf{y}_{\text{obs}} - \boldsymbol{\varphi}_{\text{obs}})^\top \Omega_{o|q}^{-1} (\mathbf{y}_{\text{obs}} - \boldsymbol{\varphi}_{\text{obs}})}{2} \right\} \\ \tilde{p}(q_{km} | \nu_m) &= q_{km}^{(\nu_m/2)-1} \exp \left\{ -\frac{\nu_m q_{km}}{2} \right\}.\end{aligned}\quad (10)$$

3.3. When ν_m is unknown

Both the E-step and M-step becomes a bit more involved when ν_m also needs to be estimated. For such cases, we can employ the ECME (expectation-conditional maximization either) algorithm [6, 8, 7] —an extension of EM algorithm—to speed up the computation. Let $\boldsymbol{\theta} = \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2\}$, where $\boldsymbol{\theta}_1 = \{\boldsymbol{\Sigma}_{wm}, \boldsymbol{\Sigma}_{dm}, \boldsymbol{\xi}_m\}$ and $\boldsymbol{\theta}_2 = \{\nu_m\}$. The ECME uses two conditional M-steps instead of one M-step:

$$\begin{aligned}\boldsymbol{\theta}_1^{[t+1]} &= \arg \max_{\boldsymbol{\theta}_1} Q(\{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2^{[t]}; \boldsymbol{\theta}^{[t]}\}) \\ \boldsymbol{\theta}_2^{[t+1]} &= \arg \max_{\boldsymbol{\theta}_2} \log p(\mathbf{y}_{\text{obs}} | \{\boldsymbol{\theta}_1^{[t+1]}, \boldsymbol{\theta}_2\}).\end{aligned}$$

Here the first CM-step is identical to the M-step described above. The second CM-step can be carried out via Newton-Raphson:

$$\nu_{m, \text{new}}^{[t+1]} = \nu_{m, \text{old}}^{[t+1]} - \frac{\ell'(\boldsymbol{\theta}_2 | \boldsymbol{\theta}_1^{[t+1]})}{\ell''(\boldsymbol{\theta}_2 | \boldsymbol{\theta}_1^{[t+1]})}, \quad (11)$$

where $\ell(\boldsymbol{\theta}_2 | \boldsymbol{\theta}_1) = \log p(\mathbf{y}_{\text{obs}} | \boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ and

$$p(\mathbf{y}_{\text{obs}} | \boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \propto \int_0^\infty \phi_{\text{obs}}(\mathbf{q}) \prod_{k,m} p(q_{km} | \nu_m) dq_{km}$$

can be evaluated numerically as before. Note here $p(q_{km} | \nu_m)$ must be the proper density of $\chi_{\nu_m}^2 / \nu_m$, not the unnormalized density as given in (10), because the normalizing constant is a function of the unknown parameter ν_m and thus it is needed in calculating the derivatives required by (11). Note also that ν_m is continuous as χ_ν^2 distribution is a special case of the Gamma distribution with the shape parameter $\alpha = \nu/2$.

4. IMPLEMENTATION, APPLICATIONS, AND RESULTS

Once $\boldsymbol{\theta}^{[t]}$ is found by our algorithm, our estimate of the true image \mathbf{f} is the posterior mean, $\mathbf{f}^{[t]} = \mathbf{W}^\top E[\mathbf{d} | \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}^{[t]}]$, where \mathbf{W} now stands for wavelet transform applied to each color component independently. This is clearly useful to a number of applications.

The algorithm is implemented using the overcomplete Daubechies 4 separable wavelets. The overcomplete expansion is accomplished by convolution (i.e. ignoring the downsample operator), and the redundant representation of wavelet coefficients are projected back to the span of the image signal by averaging together the shifted versions of inverse wavelet transforms, amounting to the model averaging in the pixel domain. The elements in \mathbf{w}_{km} are the wavelet coefficients taken from a 3×3 spatial window centered around w_{km} . Here the members of \mathbf{w}_{km} are taken only from the m -th subband, although the inclusion of parent/child wavelets is a straightforward extension. For color image \mathbf{f} , \mathbf{w}_{km} includes the wavelet coefficients from all color components belonging to their respective 3×3 spatial window centered around w_{km} . Note that there are redundant estimates of the wavelet coefficients in \mathbf{d}_{km} because of the overlapping neighborhoods of 3×3 moving window. These redundant estimates are averaged together before the inverse wavelet transform.

For our experiment, the degree of freedom ν_m is set to 0.1 for wavelet coefficients and 1 for scaling coefficients (so the second CM-step is unneeded). All other parameters are estimated by our algorithm. The initial values $\boldsymbol{\Sigma}_{dm}$ are trained using a reference image, or a typical noise-free (color) image that with no missing data. The initial noise correlation matrices $\boldsymbol{\Sigma}_{wm}$ are set to an identity matrix multiplied by the initial guess of the noise variance. The prior mean vectors $\boldsymbol{\xi}_m$ are set to 0 for the wavelet coefficients and the average color values in \mathbf{y}_{obs} for the scaling coefficients.

Because the literal execution of this algorithm is computationally intensive, we made a few approximations in our implementation (due to page constraints, the impact of such approximations will be reported in future work). First, we assume q_{km} is spatially slowly varying and therefore we effectively treat $q_{k'm} = q_{km}$ when k and k' 's belongs to the same 3×3 neighborhood (or we can view this as a model modification to (3) by assuming a common q within each neighborhood). This approximation greatly reduces the dimension of the numerical integration over q to one. Second, to greatly reduce the size of matrices in computing the regression imputation, we approximate $E[y_i | \mathbf{q}, \mathbf{y}_{\text{obs}}, \boldsymbol{\theta}]$ by $E[y_i | \mathbf{q}, \mathbf{y}_{\text{obs}_i}, \boldsymbol{\theta}]$, where $\mathbf{y}_{\text{obs}_i}$ is a subset of \mathbf{y}_{obs} restricted to a $n_i \times n_i$ neighborhood around the pixel location i ; we found a value $n_i \in (10, 20)$ typically adequate.

In the case that all noisy pixels are observed, the proposed algorithm presents a method for estimating the noise-free pixel values, \mathbf{f} , given noisy pixels \mathbf{y} . Conversely, when only a subset of the noise-free pixels are observed, the proposed algorithm offers a method to interpolate the missing pixels. Furthermore, by interpreting spectral channels as separate different *color* components of the image,




Fig. 1. Zoomed portion of the output images ($\sigma_n = 25$). Method(s) in (a) [13], (b) [13]+[1], (c) [9], (d) proposed (3rd iteration).

the interpolation method generalizes to hyper-spectral images seamlessly. In the real-world image interpolation problems, however, the observed data most certainly contain noise, sometimes to a severe degree. This is especially apparent in digital cameras, as the noise in the image sensor in a poor lighting condition is often amplified in the interpolation step known as demosaic(k)ing. While the recently published methods yield impressive results in the absence of noise, to the best of the authors' knowledge, none of the methods address the noise problems *explicitly*, with the exception of [9, 10]. Note also that the proposed algorithm readily generalizes to irregular sampling patterns (e.g. [11]) or with a choice of color filters other than red, green, and blue.

Given the page constraints, we only present experimental results on demosaicing. We assume that the image sensor, arranged in RGB Bayer pattern [12], is corrupted by white additive Gaussian noise. As the test images, we use widely available 24-bit color images often referred to as the “clown” and “lena,” and the MSE of the output images are shown in table 1 (treating all color components equally). In the presence of noise, the proposed algorithm clearly outperforms demosaicing [13], demosaicing followed by denoising [13, 1], and method in [9] that perform demosaicing and denoising simultaneously. Zoomed portion of the output images ($\sigma_n = 25$), shown in figure 1, demonstrate that the proposed algorithm reduces noise while preserving the sharpness of the image details, and introduces less noticeable artifacts when compared to the alternatives. Note that the MSE increases after several iterations partially due to overfitting by the MLE and approximate nature of our current implementation.

5. CONCLUSION

We presented a new framework for coupling the EM algorithm with the Bayesian hierarchical modeling of wavelet coefficients of image signals. The E-step in the algorithm yields an estimate of the noise-free pixel values, even if only a subset of noisy pixels are observed. The M-step finds the maximal likelihood estimates (MLE) of the model hyper-parameters. The proposed method unifies the sta-

Table 1. The MSE of the demosaiced “clown” and “lena” images corrupted by white noise with standard deviation σ_n .

method	clown		lena	
	$\sigma_n = 0$	$\sigma_n = 25$	$\sigma_n = 0$	$\sigma_n = 25$
method in [13]	50.56	492.08	26.47	578.82
methods in [13]+[1]	50.56	203.92	26.47	92.13
method in [9]	105.95	194.00	36.17	95.00
proposed, $t = 1$	45.77	163.42	21.94	83.64
proposed, $t = 2$	47.30	151.51	19.81	76.21
proposed, $t = 3$	52.54	156.48	20.26	82.43
proposed, $t = 4$	61.48	160.97	21.30	86.17
proposed, $t = 5$	70.35	165.04	22.28	89.74

tistical treatments of the missing pixels with that of the noisy observation data, thereby yielding results potentially substantially better than when interpolation and denoising are performed independently.

6. REFERENCES

- [1] J. Portilla, V. Strela, M.J. Wainwright, and E.P. Simoncelli, “Image denoising using scale mixture of Gaussians in the wavelet domain,” *IEEE Trans. Image Processing*, vol. 12, no. 11, pp. 1338–1351, November 2003.
- [2] L. Sendur and I.W. Selesnick, “Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency,” *IEEE Trans. Signal Processing*, vol. 50, 2002.
- [3] A. Pizurica, W. Philips, I. Lemahieu, and M. Acheroy, “A joint inter- and intrascale statistical model for Bayesian wavelet based image denoising,” *IEEE Trans. Image Processing*, vol. 11, no. 5, pp. 545–557, May 2002.
- [4] J. Portilla, “Full blind denoising through noise covariance estimation using Gaussian scale mixtures in the wavelet domain,” *IEEE Int’l Conf on Image Processing*, pp. 1217–1220, October 2004.
- [5] J. Portilla, “Blind non-white noise removal in images using Gaussian scale mixtures in the wavelet domain,” *Proc. of the IEEE Benelux Signal Processing Symposium*, pp. 17–20, April 2004.
- [6] G.J. McLachlan and T. Krishnan, *The EM Algorithm and Extensions*, John Wiley & Sons, Inc., 1997.
- [7] X.L. Meng and D. van Dyk, “The EM algorithm—An old folk song sung to a fast new tune (with discussions),” *J.R. Statist. Soc., Ser B*, vol. 59, pp. 511–567, 1995.
- [8] C. Liu and D.B. Rubin, “The ECME algorithm: a simple extension of EM and ECM with faster monotone convergence,” *Biometrika*, vol. 81, pp. 633–648, 1994.
- [9] K. Hirakawa and T.W. Parks, “Joint demosaicing and denoising,” *IEEE Int’l Conf on Image Processing*, September 2005.
- [10] J. Portilla, D. Otaduy, and C. Dorransoro, “Low-complexity linear demosaicing using joint spatial-chromatic image statistics,” *IEEE Int’l Conf on Image Processing*, September 2005.
- [11] M. Parmar and S.J. Reeves, “A perceptually based design methodology for color filter arrays,” *IEEE Int’l Conf on Acoustics, Speech, and Signal Processing*, 2004.
- [12] B.E. Bayer, “Color imaging array,” US Patent 3 971 065, 1976.
- [13] B.K. Gunturk, Y. Altunbasak, and R.M. Mersereau, “Color plane interpolation using alternating projections,” *IEEE Trans. Image Processing*, vol. 11, no. 9, pp. 997–1013, September 2002.