



Heriot-Watt University  
Research Gateway

## Patch-Based Image Restoration using Expectation Propagation

**Citation for published version:**

Yao, D, McLaughlin, S & Altmann, Y 2022, 'Patch-Based Image Restoration using Expectation Propagation', *SIAM Journal on Imaging Sciences*, vol. 15, no. 1, pp. 192-227. <https://doi.org/10.1137/21M1427541>

**Digital Object Identifier (DOI):**

[10.1137/21M1427541](https://doi.org/10.1137/21M1427541)

**Link:**

[Link to publication record in Heriot-Watt Research Portal](#)

**Document Version:**

Publisher's PDF, also known as Version of record

**Published In:**

SIAM Journal on Imaging Sciences

**Publisher Rights Statement:**

© 2022 SIAM.

**General rights**

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [open.access@hw.ac.uk](mailto:open.access@hw.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.

## Patch-Based Image Restoration Using Expectation Propagation\*

Dan Yao<sup>†</sup>, Stephen McLaughlin<sup>†</sup>, and Yoann Altmann<sup>†</sup>

**Abstract.** This paper presents a new Expectation Propagation (EP) framework for image restoration using patch-based prior distributions. While Monte Carlo techniques are classically used to sample from intractable posterior distributions, they can suffer from scalability issues in high-dimensional inference problems such as image restoration. To address this issue, EP is used here to approximate the posterior distributions using products of multivariate Gaussian densities. Moreover, imposing structural constraints on the covariance matrices of these densities allows for greater scalability and distributed computation. While the method is naturally suited to handle additive Gaussian observation noise, it can also be extended to non-Gaussian noise. Experiments conducted for denoising, inpainting, and deconvolution problems with Gaussian and Poisson noise illustrate the potential benefits of such a flexible approximate Bayesian method for uncertainty quantification in imaging problems, at a reduced computational cost compared to sampling techniques.

**Key words.** expectation propagation, patch-based prior, structured covariance matrices, Gaussian mixture model, uncertainty quantification, image restoration

**AMS subject classifications.** 62E17, 62F15, 62F30, 62H10, 68W25

**DOI.** 10.1137/21M1427541

### 1. Introduction.

**1.1. Problem formulation.** In this paper, we address the problem of image restoration, which consists of estimating an unknown image from its degraded observation, e.g., noisy, blurry, or missing pixels. The (vectorized) observed image  $\mathbf{y} \in \mathbb{R}^N$  is modeled as a known linear transformation of the unknown image of interest  $\mathbf{x} \in \mathbb{R}^N$ . The linear transformation is denoted by  $\mathbf{H}\mathbf{x}$ , and the matrix  $\mathbf{H} \in \mathbb{R}^{N \times N}$  is either diagonal (e.g., as in denoising and inpainting problems) or diagonalizable in a transformed domain (e.g., as in deconvolution problems). The observation noise can be either Gaussian, independent and identically distributed (i.i.d.), or more complex, e.g., non-Gaussian. Adopting a Bayesian approach, the observation model described by the likelihood function  $f_{\mathbf{y}|\mathbf{x}}(\mathbf{y}|\mathbf{H}\mathbf{x})$ , and a prior distribution  $f_{\mathbf{x}}(\mathbf{x}|\boldsymbol{\theta})$  parameterized by  $\boldsymbol{\theta}$  and based on image patches, are adopted to perform image restoration. Following Bayes'

\*Received by the editors June 17, 2021; accepted for publication (in revised form) November 3, 2021; published electronically February 15, 2022.

<https://doi.org/10.1137/21M1427541>

**Funding:** The work of the authors was supported by the UK Royal Academy of Engineering under the Research Fellowship Scheme (RF201617/16/31), the Engineering and Physical Sciences Research Council of the UK (EPSRC) grant EP/S000631/1, and the UK MOD University Defence Research Collaboration (UDRC) in Signal Processing.

<sup>†</sup>School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, EH14 4AS, Scotland, UK ([dy14@hw.ac.uk](mailto:dy14@hw.ac.uk), [S.McLaughlin@hw.ac.uk](mailto:S.McLaughlin@hw.ac.uk), [Y.Altmann@hw.ac.uk](mailto:Y.Altmann@hw.ac.uk)).

theorem, the posterior distribution of  $\mathbf{x}$  is

$$(1.1) \quad f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) = \frac{f_{\mathbf{y}|\mathbf{x}}(\mathbf{y}|\mathbf{H}\mathbf{x})f_{\mathbf{x}}(\mathbf{x}|\boldsymbol{\theta})}{\int f_{\mathbf{y}|\mathbf{x}}(\mathbf{y}|\mathbf{H}\mathbf{x})f_{\mathbf{x}}(\mathbf{x}|\boldsymbol{\theta})d\mathbf{x}}.$$

Exact inference beyond maximum a posteriori (MAP) estimation using (1.1) is however extremely difficult for high-dimensional imaging problems, due to the high-dimensional, usually intractable, integral in the denominator. Prior to describing the proposed approximate Bayesian method enabling scalable inference, we briefly describe related approaches previously investigated for Bayesian image restoration.

## 1.2. Related work.

*Designing the prior model.* The choice of suitable prior models depends on the amount of information available about the images expected to be recovered and generally relies on a trade-off between the quality of the resulting images and the computational complexity of the algorithm(s) required to perform Bayesian inference. Traditional image priors rely on explicit distributions modeling statistical properties of images in some transform domain [1, 2, 3], or Markov random fields (MRFs) [4]. Such prior models are often generic, i.e., they do not encode strong prior information and have been successfully applied to a variety of imaging problems, e.g., medical imaging and astroimaging, for the restoration of images corrupted by different noise models.

Over the last few years, patch-based image restoration methods have also received a lot of interest. Such approaches are particularly attractive from a computational point of view, as learning patch-based priors and using them to restore patches/images can be simpler than learning and using global image priors [5, 6]. In practice, overlapping patches are often found to be crucial to achieve performance close to that of state-of-the-art image restoration methods [7]. There are two main categories of methods in the literature to learn/build patch-based priors. In the first category, patch-based prior models are learned on image patches which are assumed to be independent. Such methods include, for example, [5, 8, 9, 10, 11] using Gaussian mixture models (GMMs) as a patch-based prior. Conversely, the second category of methods accounts for dependencies between overlapping patches while learning a patch-based prior. Fields-of-Experts (FoE) priors [12] are typical examples of such methods. However, training FoE models is computationally intensive compared to the first category of methods. In this paper, to keep the overall inference process (prior design and image restoration) computationally tractable, we consider expressive GMMs built using tools from the first category of methods, as priors of image patches. To alleviate the usual blocky artifacts when restoring images from non-overlapping patch estimates, here we use GMM priors for different partitions of the image of interest. Individual image posteriors are obtained from the different partitions, and they are combined to obtain the final image estimates (approximate means and marginal variances).

Recently, more advanced priors, such as plug-and-play priors [13, 14] and priors trained by a deep neural network [15, 16, 17, 18], have become popular in image restoration methods and have shown leading performance in many image restoration tasks. However, most of these priors are limited to delivering point estimates, without clear rules to quantify the resulting

uncertainties. The GMM patch-based prior considered in this work helps to build a flexible and scalable inference framework to provide both point estimates and pixel-wise posterior uncertainty measures. Since this prior models non-overlapping image patches independently, it allows for greater scalability in solving high-dimensional imaging problems. Being a conjugate prior for a Gaussian observation model [19], it is extended in this work to solve more complex image restoration problems beyond denoising with Gaussian noise, such as observations corrupted by Poisson noise.

*Choosing appropriate point estimators and summary statistics.* The majority of Bayesian methods using patch-based image priors adopt the MAP estimator as the preferred point estimator of  $\mathbf{x}$  as it can be computed using powerful optimization tools. In the case of patch-based priors, e.g., [5, 9, 10, 11, 12], the resulting posterior distribution may not be convex or unimodal. Nonetheless, convex optimization tools can be used to optimize (local) bounds of the log-posterior distribution [9]. An alternative point estimator is the minimum mean squared error (MMSE) estimator, which, when combined with associated uncertainty measurement, is particularly useful in problems where the posterior distributions are skewed. This is typically the case in imaging situations where the observations are corrupted by Poisson-like, multiplicative noise. Unfortunately, exact MMSE inference for high-dimensional image restoration problems is often intractable or computationally expensive. In the context of patch-based priors, while MMSE estimation is not generally possible at the image level, it is however possible to compute MMSE patch estimates and to combine these estimates by averaging. The MMSE estimates can be either obtained analytically [8] or approximated by Monte Carlo sampling [6]. These approaches, which have been proposed for image denoising and inpainting, are closely related to our work. However, we extend their capabilities to image deconvolution, i.e., when degradation operator  $\mathbf{H}$  is non-diagonal.

While most MAP-based methods can produce satisfactory point estimates for restored images, the algorithms are generally unable to directly quantify the uncertainty in the solutions delivered [20]. For log-concave posteriors, however, MAP estimation methods based on convex optimization can identify regions of high-posterior densities that can be used in hypothesis testing [21]. Although numerous patch-based image restoration methods have been proposed (e.g., [5, 6, 8, 9, 10, 12]), such priors are rarely studied beyond MAP estimation. This is probably because several challenges arise when considering MMSE estimation and uncertainty quantification at the image level using priors based on overlapping patches.

*Approximate estimation.* When exact computation of posteriori summary statistics is not possible, approximation strategies have been proposed to approximate the intractable posterior distribution. These strategies can be stochastic, deterministic, or hybrid. Stochastic methods are carried out by sampling, such as importance sampling [6] or Markov chain Monte Carlo (MCMC) sampling. However, sampling methods for accurate approximation of high-dimensional integrals are still computationally expensive and not yet fully scalable for fast inference. While estimating posterior means might only require relatively short Markov chains, estimating higher-order moments is much more complicated, especially when the samplers used suffer from poor mixing properties as the dimension of the problem increases. Recent advances in high-dimensional samplers for non-smooth posterior distributions, e.g., using proximal MCMC methods [22], have led to significant speed-up, at the cost of a tunable accuracy loss induced by model augmentation and/or posterior smoothing. Variational inference repre-

sents another family of approximate methods which complement sampling methods [23, 24]. Their basic principle is to approximate the target distribution  $f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$  by a more tractable distribution whose moments are easier to compute. This family includes Variational Bayes (VB) methods and methods relying on Expectation Propagation (EP) [25, 26, 27], which has recently received growing interest for high-dimensional problems. The method proposed in this work relies on EP, whose motivations and imaging applications will be further reviewed in the next paragraph.

*Expectation Propagation for imaging problems.* EP algorithms have been successfully applied to imaging problems with different observation models and different image priors. For example, in [28], a Laplace prior enforcing sparsity of the wavelet coefficients or of the image gradient was adopted to address image deconvolution problems in the presence of Gaussian noise. Gaussian processes and sparse linear models [29, 30] were considered as prior distributions when combined with a Poisson observation model. Despite the success of EP when applied to (generalized) linear regression models [31, 32, 33, 34] and imaging restoration problems [35, 36, 37, 38], to the best of our knowledge, EP has not yet been combined with patch-based priors.

Identifying uncertainties and posterior correlation in high-dimensional image restoration problems usually requires the computation of large covariance/precision matrices. This computational cost is presented when performing exact inference but also when using approximate methods like VB/EP. Two main approximate covariance matrix types have been exploited in EP methods to model, or not, the correlation between variables. Using full covariance matrices requires either the inversion of large matrices or a large number of sequential updates leveraging rank-1 optimization of large matrices [29, 38, 39]. To reduce computational time and memory requirements, diagonal covariance matrices are often used within EP. This effectively neglects local correlation between groups of variables. However, it has been shown that if significant correlation is omitted, EP can produce poor results [40].

This paper considers block-diagonal covariance matrices to capture partial correlation between variables. Instead of considering a vanilla EP factorization and implementation resulting in a large number of sequential updates, the proposed EP method consists of only two or three factors, depending on the noise model considered. This leads to EP algorithms with only two to three sequential updates at each iteration, each update being highly parallelizable and allowing distributed inference. This is a strategy similar to the factorization scheme used in AMP/GAMP/VAMP [41, 42, 43], which enables faster inference. However, in contrast to AMP/GAMP/VAMP, not all the covariance matrices are constrained to be isotropic or diagonal in the proposed EP method.

The goal of this work is to propose a new framework of EP algorithms with existing patch-based priors, enabling approximate uncertainty quantification and visualization for imaging problems. The main contributions of this paper can be summarized as follows.

### 1.3. Contributions.

1. A new EP algorithm with a GMM patch-based prior is, to the best of our knowledge, proposed for the first time for image restoration. Approximate MMSE estimates and associated posterior uncertainties are directly obtained from the moments of the Gaussian approximation of the exact posterior distribution. The algorithm can be

applied to large images thanks to the structure imposed to the covariance matrices of the approximating distributions.

2. A variational Expectation Maximization (EM) algorithm is also used to estimate GMM hyperparameters using the EP approximation of the exact posterior distribution. This allows the prior model to be adjusted for each image, and the resulting algorithm does not require significant user supervision.
3. While the method is naturally suited to handle Gaussian noise models, it is extended for non-Gaussian noise cases, and the Poisson noise model is used as an illustration.

The remainder of this paper is organized as follows. A new patch-based expectation propagation algorithm for Gaussian observation models is proposed in section 2. In section 3, the algorithm is generalized to non-Gaussian observation models, using Poisson noise as a running example. Section 4 describes the final image restoration strategy relying on model averaging. In section 5, an EP-EM strategy is proposed to estimate hyperparameters of the GMM priors. Experimental results are presented in section 6. Finally, conclusions and discussions are reported in section 7.

## 2. Patch-based EP algorithm with Gaussian noise.

**2.1. Exact Bayesian model with GMM prior.** Given a vectorized image  $\mathbf{x}$  composed of  $\sqrt{N} \times \sqrt{N}$  pixels, we consider a set of  $J$  small non-overlapping patches  $\{\mathbf{x}_j\}_{j=1,\dots,J}$  of size  $\sqrt{r} \times \sqrt{r}$ . Each pixel in the image is assigned to a single patch, leading to  $J = N/r$ . The patches  $\{\mathbf{x}_j\}_{j=1,\dots,J}$  are independently assigned the same GMM prior model

$$(2.1) \quad \mathbf{x}_j | \boldsymbol{\theta} \sim \sum_{k=1}^K \omega_k \mathcal{N}(\mathbf{x}_j; \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k),$$

where  $K$  is the number of mixture components,  $\omega_k$ ,  $\tilde{\boldsymbol{\mu}}_k$ ,  $\tilde{\mathbf{C}}_k$  are respectively the weight, mean, and covariance matrix of the  $k$ th Gaussian component, and  $\boldsymbol{\theta} = \{\omega_k, \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k\}_{k=1,\dots,K}$ . Moreover, the positive weights  $\{\omega_k\}_{k=1,\dots,K}$  sum to one. In this section and sections 3 and 4, the GMM parameters in  $\boldsymbol{\theta}$  are assumed to be known. Possible procedures to learn these parameters from training images will be discussed in section 5, together with our proposed method to refine them. The global image prior model  $f_x(\mathbf{x} | \boldsymbol{\theta})$  can be expressed as

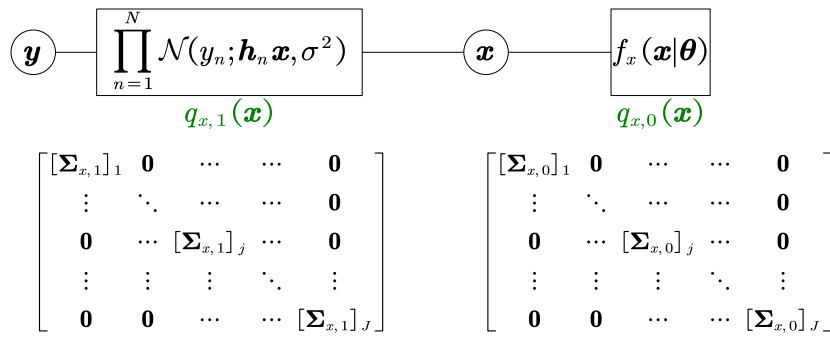
$$(2.2) \quad f_x(\mathbf{x} | \boldsymbol{\theta}) = \prod_{j=1}^J \sum_{k=1}^K \omega_k \mathcal{N}(\mathbf{x}_j; \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k).$$

In this section, we assume that the observations  $\mathbf{y}$  are corrupted by i.i.d. Gaussian noise with known noise variance  $\sigma^2$ . Thus, the likelihood function can be expressed as

$$(2.3) \quad f_{y|x}(\mathbf{y} | \mathbf{H}\mathbf{x}) = \prod_{n=1}^N \mathcal{N}(y_n; \mathbf{h}_n \mathbf{x}, \sigma^2),$$

where  $\{\mathbf{h}_n\}_{n=1,\dots,N} \in \mathbb{R}^{1 \times N}$  are the row vectors of  $\mathbf{H}$ . The resulting exact posterior distribution of  $\mathbf{x}$  is given, up to a multiplicative constant, by

$$(2.4) \quad f(\mathbf{x} | \mathbf{y}, \boldsymbol{\theta}) \propto \left[ \prod_{n=1}^N \mathcal{N}(y_n; \mathbf{h}_n \mathbf{x}, \sigma^2) \right] f_x(\mathbf{x} | \boldsymbol{\theta}).$$



**Figure 1.** Factor graph used to perform EP approximation for Gaussian likelihood and GMM patch-based prior. The rectangular boxes (resp., circles) represent the factors (resp., variables) involved in the factorized exact posterior distribution. The corresponding EP approximating factors are shown in green. The covariance matrices  $\Sigma_{x,1}$ ,  $\Sigma_{x,0}$  of the EP approximating factors are block-diagonal.

The goal of the proposed EP algorithm is to find a Gaussian distribution  $Q(\mathbf{x})$  that closely approximates the exact posterior  $f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$ ; in particular, when computing the MMSE estimate  $\mathbb{E}_{f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})}[\mathbf{x}]$  is intractable. Using the tractable approximating distribution  $Q(\mathbf{x})$ , the exact MMSE estimate can be approximated by  $\mathbb{E}_{f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})}[\mathbf{x}] \approx \mathbb{E}_{Q(\mathbf{x})}[\mathbf{x}]$ , and the posterior covariance matrix can be approximated by the covariance matrix of  $Q(\mathbf{x})$ .

Note that in cases where  $\mathbf{H}^T \mathbf{H}$  is diagonal, e.g., in denoising or inpainting problems, the EP algorithm discussed in this section is not required to perform a posteriori inference. Indeed, as discussed in [8], in that case, the posterior distribution in (2.4) is tractable and reduces to patch-wise independent GMMs. However, we still discuss this case as a starting point of this paper, as the resulting EP updates can be used for more complex problems where the noise is non-Gaussian, as in section 3.

**2.2. EP approximation strategy.** As alluded to above, EP methods rely on approximating a distribution by a product of (unnormalized) probability density function (p.d.f.) or probability mass function (p.m.f.) from the exponential family. The factors of that product are then updated sequentially. To alleviate a large number of sequential updates, we consider an approximation of  $f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$  with two factors. The first factor, which relates to  $f_x(\mathbf{x}|\boldsymbol{\theta})$  (seen as a function of  $\mathbf{x}$ ) and which depends implicitly on  $\boldsymbol{\theta}$ , is denoted by  $q_{x,0}(\mathbf{x})$ . The second factor, which relates to  $f_{y|x}(\mathbf{y}|\mathbf{H}\mathbf{x})$  (seen as a function of  $\mathbf{x}$ ) and which implicitly depends on  $\mathbf{y}$ , is denoted by  $q_{x,1}(\mathbf{x})$ . These two factors are, for simplicity, unnormalized multivariate Gaussian densities whose parameters are defined as

$$(2.5) \quad q_{x,0}(\mathbf{x}) \propto \mathcal{N}(\mathbf{x}; \mathbf{m}_{x,0}, \Sigma_{x,0}) \quad \text{and} \quad q_{x,1}(\mathbf{x}) \propto \mathcal{N}(\mathbf{x}; \mathbf{m}_{x,1}, \Sigma_{x,1}).$$

A factor graph representing this EP factorization is depicted in Figure 1. The resulting approximation of  $f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$  is given by  $Q(\mathbf{x}) \propto q_{x,1}(\mathbf{x})q_{x,0}(\mathbf{x})$ . To find the variational parameters of  $Q(\mathbf{x})$ , EP reduces to solving iteratively the following two KL divergence minimization

problems:

$$(2.6a) \quad \left\{ \begin{array}{l} \min_{q_{x,0}(\mathbf{x})} KL(f_x(\mathbf{x}|\boldsymbol{\theta})q_{x,1}(\mathbf{x})||q_{x,0}(\mathbf{x})q_{x,1}(\mathbf{x})), \\ \min_{q_{x,1}(\mathbf{x})} KL\left(\prod_{n=1}^N \mathcal{N}(y_n; \mathbf{h}_n \mathbf{x}, \sigma^2)q_{x,0}(\mathbf{x})||q_{x,1}(\mathbf{x})q_{x,0}(\mathbf{x})\right). \end{array} \right.$$

Traditionally, the densities obtained by removing a single factor from  $Q(\mathbf{x})$ , i.e.,  $q_{x,\setminus i}(\mathbf{x}) \propto Q(\mathbf{x})/q_{x,i}(\mathbf{x})$  ( $\forall i \in \{0,1\}$ ), are referred to as cavity distributions. Here, we simply have  $q_{x,\setminus 0}(\mathbf{x}) = q_{x,1}(\mathbf{x})$  and  $q_{x,\setminus 1}(\mathbf{x}) = q_{x,0}(\mathbf{x})$ . Moreover, the (unnormalized) densities obtained by multiplying these cavity distributions by their corresponding true factors are referred to as tilted distributions. Here, we have two tilted distributions,  $P_0(\mathbf{x}) \propto f_x(\mathbf{x}|\boldsymbol{\theta})q_{x,1}(\mathbf{x})$  and  $P_1(\mathbf{x}) \propto \prod_{n=1}^N \mathcal{N}(y_n; \mathbf{h}_n \mathbf{x}, \sigma^2)q_{x,0}(\mathbf{x})$ , each being associated with one of the KL divergence minimization problems in (2.6).

Within the traditional EP framework where the problems in (2.6) are solved without additional constraints on the covariance matrices  $\boldsymbol{\Sigma}_{x,0}$  and  $\boldsymbol{\Sigma}_{x,1}$ , solving the KL divergence minimization problems reduces to (i) computing the mean and covariance matrix of each tilted distribution, denoted by  $\mathbb{E}_{P_i}[\mathbf{x}]$  and  $\text{Cov}_{P_i}(\mathbf{x})$  (for  $i \in \{0,1\}$ ), and (ii) matching those moments with those of  $Q(\mathbf{x})$ , denoted by  $\mathbf{m}_*$  and  $\boldsymbol{\Sigma}_*$ . Since  $Q(\mathbf{x})$  is the product of two Gaussian densities, these approximate moments are given by

$$(2.7) \quad \boldsymbol{\Sigma}_*^{-1} = \boldsymbol{\Sigma}_{x,0}^{-1} + \boldsymbol{\Sigma}_{x,1}^{-1}, \quad \mathbf{m}_* = \boldsymbol{\Sigma}_* \left( \boldsymbol{\Sigma}_{x,0}^{-1} \mathbf{m}_{x,0} + \boldsymbol{\Sigma}_{x,1}^{-1} \mathbf{m}_{x,1} \right).$$

At each EP iteration, the algorithm produces, sequentially, new values of  $(\mathbf{m}_{x,0}, \boldsymbol{\Sigma}_{x,0})$  and  $(\mathbf{m}_{x,1}, \boldsymbol{\Sigma}_{x,1})$  that minimize the first and second lines of (2.6), respectively. However, this simple strategy often does not apply directly in the multivariate case as, for instance, the covariance matrix  $\boldsymbol{\Sigma}_{x,1}^{-1} = (\text{Cov}_{P_1}(\mathbf{x}))^{-1} - \boldsymbol{\Sigma}_{x,0}^{-1}$  is not ensured to be positive definite. To alleviate this issue, the covariance matrices  $\boldsymbol{\Sigma}_{x,0}$  and  $\boldsymbol{\Sigma}_{x,1}$  are often constrained to be diagonal or isotropic. This drastically simplifies the resolution of the problems in (2.6) when subject to the symmetry and positive-definiteness constraints applied to the covariance matrices. However, while imposing such diagonal structural constraints on  $\boldsymbol{\Sigma}_{x,0}$  and  $\boldsymbol{\Sigma}_{x,1}$  can stabilize the EP algorithm, it constrains  $Q(\mathbf{x})$  to neglect posterior correlations between the elements of  $\mathbf{x}$  and thus potentially degrades the quality of the approximation (as will be further discussed in section 6). In this work, we propose to use less restrictive constraints and allow  $\boldsymbol{\Sigma}_{x,0}$  and  $\boldsymbol{\Sigma}_{x,1}$  to be block-diagonal matrices, as discussed next.

**2.3. Block-diagonal covariance matrix approximation.** Spatial correlation between pixels is induced by the GMM priors (a priori known correlation), but it can also be induced by the likelihood factor, e.g., in deconvolution problems. While estimating and handling full, unstructured covariance matrices is not tractable for large images due to computationally intensive matrix inversions, it is possible to partially capture posterior correlation by imposing structural constraints on  $(\boldsymbol{\Sigma}_{x,1}, \boldsymbol{\Sigma}_{x,0})$ , and thus on  $\boldsymbol{\Sigma}_*$ . Enforcing block-diagonal structures appears to be a natural and computationally attractive option. Indeed, the prior model in (2.2) induces, by construction, spatial correlation within image patches and a block-diagonal structure in the prior covariance matrix of  $\mathbf{x}$  (provided that the pixels are sorted properly in



$\mathbf{x}$ ). Moreover, the degradation operator  $\mathbf{H}$  formed by convolution kernels is diagonally dominant. Thus, we propose to use the same block-diagonal structure for  $\Sigma_{x,1}$  and  $\Sigma_{x,0}$ , which is the  $r \times r$  block structure induced by the patch-based prior. For any patch-based image partition, it is possible to use this block-diagonal structure by re-ordering the observations in  $\mathbf{y}$ , the pixels in  $\mathbf{x}$ , and the rows of  $\mathbf{H}$ .

**2.4. EP update schemes with block-diagonal covariance matrices.** We can now discuss how the problems in (2.6) can be solved efficiently assuming that  $\Sigma_{x,1}$  and  $\Sigma_{x,0}$  share the same block-diagonal structure. The simplest is the update in (2.6a).

*EP update of  $q_{x,0}(\mathbf{x})$ .* The first step of any EP update based on Gaussian approximations consists of computing the mean and covariance matrix of the tilted distribution. Since  $f_x(\mathbf{x}|\boldsymbol{\theta})$  forms a product of independent GMMs and since  $q_{x,1}(\mathbf{x})$  is a multivariate Gaussian density with a block-diagonal covariance matrix matching the covariance structure of  $f_x(\mathbf{x}|\boldsymbol{\theta})$ , the tilted distribution  $P_0(\mathbf{x}) = f_x(\mathbf{x}|\boldsymbol{\theta})q_{x,1}(\mathbf{x})$  is also a product of  $J$  independent GMMs, whose parameters can be computed analytically and independently. For the  $j$ th image patch, the mean  $\mathbb{E}_{P_0}[\mathbf{x}_j]$  and covariance matrix  $\text{Cov}_{P_0}(\mathbf{x}_j)$  are given by

$$(2.8) \quad \mathbb{E}_{P_0}[\mathbf{x}_j] = \sum_{k=1}^K \hat{\omega}_{j,k} \hat{\boldsymbol{\mu}}_{j,k}, \quad \text{Cov}_{P_0}(\mathbf{x}_j) = \sum_{k=1}^K \hat{\omega}_{j,k} (\hat{\boldsymbol{\mu}}_{j,k} \hat{\boldsymbol{\mu}}_{j,k}^T + \hat{\mathbf{C}}_{j,k}) - \mathbb{E}_{P_0}[\mathbf{x}_j] \mathbb{E}_{P_0}^T[\mathbf{x}_j],$$

where  $\hat{\omega}_{j,k}$ ,  $\hat{\boldsymbol{\mu}}_{j,k}$ , and  $\hat{\mathbf{C}}_{j,k}$  are obtained by

$$(2.9) \quad \begin{cases} \hat{\omega}_{j,k} = \frac{\omega_k \mathcal{N}([\mathbf{m}_{x,1}]_j; \tilde{\boldsymbol{\mu}}_k, [\Sigma_{x,1}]_j + \tilde{\mathbf{C}}_k)}{\sum_{k=1}^K \omega_k \mathcal{N}([\mathbf{m}_{x,1}]_j; \tilde{\boldsymbol{\mu}}_k, [\Sigma_{x,1}]_j + \tilde{\mathbf{C}}_k)}, \\ \hat{\mathbf{C}}_{j,k} = \left( [\Sigma_{x,1}]_j^{-1} + \tilde{\mathbf{C}}_k^{-1} \right)^{-1}, \\ \hat{\boldsymbol{\mu}}_{j,k} = \hat{\mathbf{C}}_k \left( [\Sigma_{x,1}]_j^{-1} [\mathbf{m}_{x,1}]_j + \tilde{\mathbf{C}}_k^{-1} \tilde{\boldsymbol{\mu}}_k \right), \end{cases}$$

where  $[\Sigma_{x,1}]_j$  denotes the block on the diagonal of  $\Sigma_{x,1}$  associated with the marginal covariance matrix of the  $j$ th patch and  $[\mathbf{m}_{x,1}]_j$  corresponds to the mean of that patch (extracted from  $\mathbf{m}_{x,1}$ ). Once the mean and covariance matrix of  $P_0(\mathbf{x})$  are computed, the next step consists of the actual update of  $(\mathbf{m}_{x,0}, \Sigma_{x,0})$ . Again, this can be done effectively by leveraging the block structure of the covariance matrices, as the blocks can be processed independently. Let  $\Omega_0 = \Sigma_{x,0}^{-1}$ ,  $\Omega_1 = \Sigma_{x,1}^{-1}$  be the precision matrices of  $q_{x,0}(\mathbf{x})$  and  $q_{x,1}(\mathbf{x})$ , respectively. These matrices can be easily obtained by  $J$  inversions of small matrices of size  $r \times r$ ,  $\Omega_0 = \text{diag}([\Sigma_{x,0}]_1^{-1}, \dots, [\Sigma_{x,0}]_J^{-1})$ ,  $\Omega_1 = \text{diag}([\Sigma_{x,1}]_1^{-1}, \dots, [\Sigma_{x,1}]_J^{-1})$ . The KL divergence in (2.6a) can be expressed as

$$(2.10) \quad KL(P_0(\mathbf{x})||Q(\mathbf{x})) = \int P_0(\mathbf{x}) \left[ -\frac{1}{2} \log(\det(\Omega_*)) + \frac{1}{2} (\mathbf{x} - \mathbf{m}_*)^T \Omega_* (\mathbf{x} - \mathbf{m}_*) \right] d\mathbf{x} + \kappa,$$

where  $\Omega_* = \Omega_0 + \Omega_1$  and  $\kappa = \int P_0(\mathbf{x}) \log P_0(\mathbf{x}) d\mathbf{x} + \frac{N}{2} \log(2\pi)$  is a constant independent of  $Q(\mathbf{x})$ . Minimizing (2.10) is equivalent to minimizing the loss function  $F_0(\mathbf{m}_*, \Omega_*)$  defined as

$$(2.11) \quad F_0(\mathbf{m}_*, \Omega_*) = -\log(\det(\Omega_*)) + \langle \Omega_*, \text{Cov}_{P_0}(\mathbf{x}) \rangle + (\mathbb{E}_{P_0}[\mathbf{x}] - \mathbf{m}_*)^T \Omega_* (\mathbb{E}_{P_0}[\mathbf{x}] - \mathbf{m}_*),$$

with  $\langle \mathbf{\Omega}_*, \text{Cov}_{P_0}(\mathbf{x}) \rangle = \text{Trace}(\mathbf{\Omega}_* \text{Cov}_{P_0}(\mathbf{x}))$ .  $F_0(\mathbf{m}_*, \mathbf{\Omega}_*)$  is convex with respect to (w.r.t.)  $\mathbf{m}_*$  and minimized when  $\mathbf{m}_* = \mathbb{E}_{P_0}[\mathbf{x}]$  (no constraints are applied to  $\mathbf{m}_*$ ). In that case, the loss function (as a function of  $\mathbf{\Omega}_*$ ) reduces to

$$(2.12) \quad F_1(\mathbf{\Omega}_*) = -(\log(\det \mathbf{\Omega}_*)) + \langle \mathbf{\Omega}_*, \text{Cov}_{P_0}(\mathbf{x}) \rangle.$$

By taking advantage of the block-diagonal structure of the covariance matrix  $\text{Cov}_{P_0}(\mathbf{x})$  of the tilted distribution and of  $\mathbf{\Sigma}_{x,1}$ ,  $F_1(\mathbf{\Omega}_*)$  can be re-written as

$$(2.13) \quad F_1(\mathbf{\Omega}_0) \propto \sum_{j=1}^J -\log(\det(\mathbf{\Omega}_{0,j} + \mathbf{\Omega}_{1,j})) + \langle (\mathbf{\Omega}_{0,j} + \mathbf{\Omega}_{1,j}), \text{Cov}_{P_0}(\mathbf{x}_j) \rangle,$$

which can be minimized patch-wise w.r.t.  $\{\mathbf{\Omega}_{0,j}\}_{j=1,\dots,J}$  via gradient-based methods. Here, we use an iterative method, where at the  $(t)$ th iteration we have

$$(2.14) \quad \mathbf{\Omega}_{0,j}^{(t)} = \mathbf{\Omega}_{0,j}^{(t-1)} - \lambda^{(t)} [\text{Cov}_{P_0}(\mathbf{x}_j) - (\mathbf{\Omega}_{0,j}^{(t-1)} + \mathbf{\Omega}_{1,j})^{-1}].$$

At each iteration, the initial step-size is chosen by the Barzilai–Borwein method [44], and it is then adjusted by backtracking line search to make sure the KL divergence decreases between successive iterations and that  $\mathbf{\Omega}_{0,j}^{(t)} \in \mathcal{S}_{++}^r$ , where  $\mathcal{S}_{++}^r$  denotes the space of  $r \times r$  symmetric, positive definite matrices. Once  $\mathbf{\Omega}_{0,j}$  is obtained, the  $j$ th diagonal block of  $\mathbf{\Sigma}_{x,0}$  and  $j$ th subset of  $\mathbf{m}_{x,0}$  are updated using

$$(2.15) \quad [\mathbf{\Sigma}_{0,x}]_j = \mathbf{\Omega}_{0,j}^{-1}, \quad [\mathbf{m}_{x,0}]_j = [\mathbf{\Sigma}_{0,x}]_j ((\mathbf{\Omega}_{0,j} + \mathbf{\Omega}_{1,j}) \mathbb{E}_{P_0}[\mathbf{x}_j] - \mathbf{\Omega}_{1,j} [\mathbf{m}_{x,1}]_j).$$

*EP update of  $q_{x,1}(\mathbf{x})$ .* In (2.6b), the tilted distribution  $P_1(\mathbf{x}) \propto \prod_{n=1}^N \mathcal{N}(y_n; \mathbf{h}_n \mathbf{x}, \sigma^2) q_{x,0}(\mathbf{x})$  is the product of two Gaussian densities and is thus also a multivariate Gaussian density. Its mean  $\mathbb{E}_{P_1}[\mathbf{x}]$  and covariance matrix  $\text{Cov}_{P_1}[\mathbf{x}]$  are given by

$$(2.16) \quad \mathbb{E}_{P_1}[\mathbf{x}] = \text{Cov}_{P_1}(\mathbf{x}) \left( \sigma^{-2} \mathbf{H}^T \mathbf{y} + \mathbf{\Sigma}_{x,0}^{-1} \mathbf{m}_{x,0} \right), \quad \text{Cov}_{P_1}(\mathbf{x}) = (\sigma^{-2} \mathbf{H}^T \mathbf{H} + \mathbf{\Sigma}_{x,0}^{-1})^{-1}.$$

The main computational bottleneck of this step is the inversion of an  $N \times N$  matrix required to compute  $\text{Cov}_{P_1}(\mathbf{x})$ , and in turn  $\mathbb{E}_{P_1}[\mathbf{x}]$ . If  $\mathbf{H}^T \mathbf{H}$  is diagonal, the inversion can be computed patch-wise using the structure of  $\mathbf{\Sigma}_{x,0}^{-1}$ . If  $\mathbf{H}^T \mathbf{H}$  is non-diagonal, the direct inversion of  $\sigma^{-2} \mathbf{H}^T \mathbf{H} + \mathbf{\Sigma}_{x,0}^{-1}$  rapidly becomes too costly. To tackle this problem, the mean of the tilted distribution  $\mathbb{E}_{P_1}[\mathbf{x}]$  is computed by solving the following problem:

$$(2.17) \quad \mathbb{E}_{P_1}[\mathbf{x}] = \underset{\mathbf{z}}{\text{argmin}} \quad \left\| \text{Cov}_{P_1}^{-1}(\mathbf{x}) \mathbf{z} - \left( \sigma^{-2} \mathbf{H}^T \mathbf{y} + \mathbf{\Sigma}_{x,0}^{-1} \mathbf{m}_{x,0} \right) \right\|_2^2,$$

since computing  $\text{Cov}_{P_1}^{-1}(\mathbf{x})$  is not as costly. This is achieved via conjugate gradient descent, which is a classical choice to solve such problems [45]. We have shown in (2.13) that only the blocks on the main diagonal of  $\text{Cov}_{P_1}(\mathbf{x})$  are required to minimize the KL divergence when the approximating distribution  $Q(\mathbf{x})$  has a block-diagonal covariance matrix. Instead of inverting  $\sigma^{-2} \mathbf{H}^T \mathbf{H} + \mathbf{\Sigma}_{x,0}^{-1}$  and selecting those blocks, we approximate those marginal

blocks using Rao–Blackwellized Monte Carlo (RBMC) [46], which provides in practice sufficiently accurate approximations. We refer interested readers to [46] for a detailed analysis about the time and memory complexity of RBMC. Once the  $J$  diagonal blocks of  $\text{Cov}_{P_1}(\mathbf{x})$ , i.e.,  $\{\text{Cov}_{P_1}(\mathbf{x}_j)\}_{j=1,\dots,J}$ , are estimated,  $(\mathbf{m}_{x,1}, \Sigma_{x,1})$  can be updated in a similar fashion to  $(\mathbf{m}_{x,0}, \Sigma_{x,0})$ . For the precision matrix of the  $j$ th block, i.e.,  $\Omega_{1,j} = [\Sigma_{x,1}]_j^{-1}$ , the  $(t)$ th gradient descent iteration reduces to

$$(2.18) \quad \Omega_{1,j}^{(t)} = \Omega_{1,j}^{(t-1)} - \lambda^{(t)}[\text{Cov}_{P_1}(\mathbf{x}_j) - (\Omega_{1,j}^{(t-1)} + \Omega_{0,j})^{-1}].$$

Once  $\Omega_{1,j}$  is obtained, the  $j$ th diagonal block of  $\Sigma_{x,1}$  and  $j$ th subset of  $\mathbf{m}_{x,1}$  are updated by

$$(2.19) \quad [\Sigma_{x,1}]_j = \Omega_{1,j}^{-1}, \quad [\mathbf{m}_{x,1}]_j = [\Sigma_{x,1}]_j ((\Omega_{1,j} + \Omega_{0,j})\mathbb{E}_{P_1}[\mathbf{x}_j] - \Omega_{0,j}[\mathbf{m}_{x,0}]_j).$$

As was done for the update of  $q_{x,0}(\mathbf{x})$ , the  $J$  blocks of  $\Sigma_{x,1}$  and  $J$  subsets of  $\mathbf{m}_{x,1}$  can be updated independently in parallel.

While  $\Sigma_{x,0}$  and  $\Sigma_{x,1}$  are constrained to be block-diagonal in (2.15) and (2.19), we can easily impose more restrictive constraints and force these matrices to be diagonal. For instance, in (2.19), if  $\Omega_{0,j}$  is non-diagonal and we want to ensure that  $\Omega_{1,j}$  is diagonal, it is sufficient to keep only the diagonal of the gradient term in (2.18), provided that the initial  $\Omega_{1,j}^{(1)}$  is diagonal. In practice, as will be shown in section 6, we have observed that constraining  $\Sigma_{x,0}$  and  $\Sigma_{x,1}$  to be diagonal leads to satisfactory results for denoising and inpainting problems where  $\mathbf{H}^T \mathbf{H}$  is diagonal, and the resulting algorithm is faster than if  $\Sigma_{x,0}$  and  $\Sigma_{x,1}$  are constrained to be only block-diagonal (see Appendix A.1 for details). Moreover, if  $\mathbf{H}^T \mathbf{H}$  is diagonal and positive-definite (e.g.,  $\mathbf{H}$  is the identity matrix), the gradient descent method in (2.18) is not needed, and we can directly set  $\Sigma_{x,1} = \sigma^2(\mathbf{H}^T \mathbf{H})^{-1}$ . Conversely, we have observed that forcing  $\Sigma_{x,0}$  and  $\Sigma_{x,1}$  to be diagonal for deconvolution problems does not lead to acceptable results in general as in this case, the resulting EP message passing structure in (2.6) does not capture or pass on important correlations originally included in the exact posterior. The pseudo-code of the proposed EP algorithm for Gaussian observation model is presented in Algorithm 1.

### 3. Patch-based EP algorithm with non-Gaussian noise.

**3.1. Poisson observation model.** For restoration of images corrupted by Poisson noise, the observations  $y_n$  for  $n = 1, \dots, N$  are classically assumed to be mutually independent, conditioned on  $\mathbf{x}$ , and the likelihood can be expressed as

$$(3.1) \quad f_{y|x}(\mathbf{y}|\mathbf{H}\mathbf{x}) = \prod_{n=1}^N \mathcal{P}_{y_n}(\mathbf{h}_n \mathbf{x}) = \prod_{n=1}^N \frac{e^{-\mathbf{h}_n \mathbf{x}} (\mathbf{h}_n \mathbf{x})^{y_n}}{y_n!},$$

where the entries of  $\{\mathbf{h}_n\}_{n=1,\dots,N}$  are non-negative. This likelihood is particularly challenging as (i) it usually requires  $\mathbf{x}$  to be positive, and (ii) it induces noise levels that change across the  $N$  observations depending on  $\mathbf{H}$  and  $\mathbf{x}$ . This makes the Poisson model an interesting case study for the quantification of posterior uncertainty.

Imposing positivity constraints on  $\mathbf{x}$  via its prior model requires modifying the GMM in (2.2). However, introducing mixtures of multivariate truncated Gaussian distributions is not a suitable option as the normalizing constants involved in the computation of the mixture

---

**Algorithm 1.** Proposed patch-based EP algorithm—Gaussian observation model.

---

**Input:**  $\mathbf{y}$ ,  $\mathbf{H}$ ,  $\sigma^2$ ,  $\{\omega_k, \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k, K\}$

**Output:**  $\boldsymbol{\Sigma}_{x,0}$ ,  $\mathbf{m}_{x,0}$ ,  $\boldsymbol{\Sigma}_{x,1}$ ,  $\mathbf{m}_{x,1}$

```

1: initialization:  $\{\mathbf{m}_{x,1}, \mathbf{m}_{x,0}\} = \mathbf{y}$ ,  $\{\boldsymbol{\Sigma}_{x,1}, \boldsymbol{\Sigma}_{x,0}\} = \sigma^2 \mathbf{I}$ 
2: while stopping criterion is not satisfied do
3:   • EP update of  $q_{x,0}(\mathbf{x})$ 
4:     compute titled mean and block covariance via (2.8)
5:     for  $j = 1, \dots, J$  in parallel do
6:       update the precision  $\boldsymbol{\Omega}_{0,j}$  of diagonal blocks
7:       update the covariance of diagonal block  $[\boldsymbol{\Sigma}_{x,0}]_j$  and subset of mean  $[\mathbf{m}_{x,0}]_j$ 
8:     end for
9:   • EP update of  $q_{x,1}(\mathbf{x})$ 
10:    compute titled mean and block covariance by RBMC in (2.16)
11:    for  $j = 1, \dots, J$  in parallel do
12:      update the precision  $\boldsymbol{\Omega}_{1,j}$  of diagonal blocks
13:      update the covariance of diagonal block  $[\boldsymbol{\Sigma}_{x,1}]_j$  and subset of mean  $[\mathbf{m}_{x,1}]_j$ 
14:    end for
15: end while

```

---

weights would be intractable. To alleviate such issues, we still consider the GMM prior model (2.2) but extend the Poisson distribution using the rectified linear transform [47]. More precisely,  $\forall u \in \mathbb{R}$ , we use the data likelihood

$$(3.2) \quad \bar{\mathcal{P}}_y(u) = \frac{u^y e^{-u}}{y!} \mathbb{I}(u > 0) + \delta(y) \mathbb{I}(u \leq 0),$$

where  $\mathbb{I}(\cdot)$  is an indicator function and  $\delta(\cdot)$  is the Dirac delta function. For  $u > 0$ ,  $\bar{\mathcal{P}}_y(u)$  reduces to the standard p.m.f. of the Poisson distribution (with positive mean), while  $y$  is forced to be 0 when  $u \leq 0$ . Combining (3.2) with the GMM patch-based prior defined in (2.2), the exact posterior is given by


$$(3.3) \quad f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) \propto \left[ \prod_{n=1}^N \bar{\mathcal{P}}_{y_n}(\mathbf{h}_n \mathbf{x}) \right] f_x(\mathbf{x}|\boldsymbol{\theta}).$$

With this exact model, if we choose the same EP factorization as in section 2,  $\mathbb{E}_{P_1}[\mathbf{x}]$  and the diagonal blocks of  $\text{Cov}_{P_1}(\mathbf{x})$  cannot be computed easily due to combination of the Poisson likelihood and the linear operator  $\mathbf{H}$ . Instead, we introduce a data augmentation scheme which fits conveniently within the EP framework and allows for simpler updates.

**3.2. Data augmentation and EP with Poisson noise.** To decouple the non-Gaussian (Poisson) noise from the linear operator  $\mathbf{H}$ , we extend the Bayesian model in (3.3) to

$$(3.4) \quad f(\mathbf{u}, \mathbf{x}|\mathbf{y}, \boldsymbol{\theta}) \propto f(\mathbf{y}|\mathbf{u}) f(\mathbf{u}|\mathbf{x}) f_x(\mathbf{x}|\boldsymbol{\theta}),$$

where  $f(\mathbf{y}|\mathbf{u}) = \prod_{n=1}^N \bar{\mathcal{P}}_{y_n}(u_n)$  and  $f(\mathbf{u}|\mathbf{x}) = \delta(\mathbf{u} - \mathbf{H}\mathbf{x}) = \prod_{n=1}^N \delta(u_n - \mathbf{h}_n \mathbf{x})$  [33]. Note that the marginal distribution obtained by integrating  $f(\mathbf{u}, \mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$  w.r.t.  $\mathbf{u}$  is the original posterior in



**Figure 2.** Factor graph used to perform EP approximation for Poisson noise model using data augmentation. The rectangular boxes (resp., circles) represent the factors (resp., variables) involved in the factorized exact posterior distribution. The corresponding EP approximating factors are shown in green. The covariance matrices  $\Sigma_{u,0}$ ,  $\Sigma_{u,1}$ ,  $\Sigma_{x,1}$ ,  $\Sigma_{x,0}$  of the EP approximating factors are diagonal, isotropic, block-diagonal, and block-diagonal, respectively.

(3.3). Instead of approximating (3.3) using  $Q(\mathbf{x})$ , we will approximate the extended posterior distribution (3.4) using a distribution  $Q(\mathbf{x}, \mathbf{u})$ , whose marginal (after integration over  $\mathbf{u}$ ) will approximate  $f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$  in (3.3).

While the model in (3.3) naturally factorizes into two factors, the extended model in (3.4) naturally factorizes using three factors due to the hierarchical structure  $\mathbf{y} - \mathbf{u} - \mathbf{x}$ . The first factor  $f(\mathbf{y}|\mathbf{u})$  only depends on  $\mathbf{u}$  (and  $\mathbf{y}$  which is observed) and will be approximated by a Gaussian density denoted  $q_{u,0}(\mathbf{u})$ . The third factor  $f_x(\mathbf{x}|\boldsymbol{\theta})$  only depends on  $\mathbf{x}$  (and  $\boldsymbol{\theta}$  which is fixed for now) and will be approximated by  $q_{x,0}(\mathbf{x})$ , as in section 2.4. The second factor  $f(\mathbf{u}|\mathbf{x})$  depends on  $\mathbf{u}$  and  $\mathbf{x}$  and is approximated by a multivariate Gaussian density  $q_1(\mathbf{u}, \mathbf{x})$  too. Although  $\mathbf{u}$  and  $\mathbf{x}$  are highly and explicitly correlated in the exact model (3.4), we use a mean-field approximation and specify  $q_1(\mathbf{u}, \mathbf{x}) = q_{u,1}(\mathbf{u})q_{x,1}(\mathbf{x})$ , where  $q_{u,1}(\cdot)$  and  $q_{x,1}(\cdot)$  are two Gaussian densities. Such approximations are common in VB methods, and here it still allows the approximated posterior means of  $\mathbf{u}$  and  $\mathbf{x}$  to be implicitly correlated. It also enables simpler EP updates, as will be shown next. The resulting factor graph and EP factorization are depicted in Figure 2, where the approximating factors are characterized by

$$(3.5) \quad q_{u,i}(\mathbf{u}) \propto \mathcal{N}(\mathbf{u}; \mathbf{m}_{u,i}, \Sigma_{u,i}) \quad \text{and} \quad q_{x,i}(\mathbf{x}) \propto \mathcal{N}(\mathbf{x}; \mathbf{m}_{x,i}, \Sigma_{x,i}) \quad \forall i \in \{0, 1\}.$$

Note that using the mean-field approximation, the approximating distribution  $Q(\mathbf{u}, \mathbf{x}) \propto q_{u,0}(\mathbf{u})q_{u,1}(\mathbf{u})q_{x,1}(\mathbf{x})q_{x,0}(\mathbf{x})$  can be written as  $Q(\mathbf{u}, \mathbf{x}) = Q(\mathbf{u})Q(\mathbf{x})$ , which gives us direct access to the moments of the marginal distribution  $\int Q(\mathbf{u}, \mathbf{x})d\mathbf{u}$  that will be used as surrogate for the intractable moments of  $f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$  in (3.3). Using these three factors, each EP iteration consists of sequentially solving the following three KL divergence minimization problems:

$$(3.6a) \quad \left\{ \begin{array}{l} \min_{q_{u,0}(\mathbf{u})} \quad KL \left( \prod_{n=1}^N \bar{\mathcal{P}}_{y_n}(u_n)q_{u,1}(\mathbf{u}) || q_{u,0}(\mathbf{u})q_{u,1}(\mathbf{u}) \right), \\ (3.6b) \quad \min_{q_{u,1}(\mathbf{u}), q_{x,1}(\mathbf{x})} \quad KL \left( \prod_{n=1}^N \delta(u_n - \mathbf{h}_n \mathbf{x})q_{u,0}(\mathbf{u})q_{x,0}(\mathbf{x}) || q_{u,1}(\mathbf{u})q_{x,1}(\mathbf{x})q_{u,0}(\mathbf{u})q_{x,0}(\mathbf{x}) \right), \\ (3.6c) \quad \min_{q_{x,0}(\mathbf{x})} \quad KL (f_x(\mathbf{x}|\boldsymbol{\theta})q_{x,1}(\mathbf{x}) || q_{x,0}(\mathbf{x})q_{x,1}(\mathbf{x}) ), \end{array} \right.$$

subject to structural constraints on different covariance matrices of the approximating factors.

**3.3. EP updates with data augmentation.** As in the Gaussian noise case, structural constraints can be enforced on the covariance matrices in the case of non-Gaussian noise with extended models. For vector  $\mathbf{x}$ , it makes sense to consider the same constraints as in the Gaussian case since  $f_x(\mathbf{x}|\boldsymbol{\theta})$  is still a patch-based prior and  $f(\mathbf{u}|\mathbf{x})$  conveys correlation between pixels via  $\mathbf{H}$ . Thus, we constrain  $\boldsymbol{\Sigma}_{x,0}$  and  $\boldsymbol{\Sigma}_{x,1}$  to be block-diagonal matrices, with the same block structure as the GMM prior. The actual likelihood term  $f(\mathbf{y}|\mathbf{u})$  does not convey correlation between the elements of  $\mathbf{u}$  since  $f(\mathbf{y}|\mathbf{u}) = \prod_{n=1}^N \bar{\mathcal{P}}_{y_n}(u_n)$ . Thus, we naturally constrain  $\boldsymbol{\Sigma}_{u,0}$  to be diagonal, and its diagonal is denoted by  $[c_{0,1}, \dots, c_{0,N}]$ . Another option would be to enforce  $\boldsymbol{\Sigma}_{u,0}$  to be isotropic, but this would not accurately capture the pixel-wise uncertainty induced by the Poisson likelihood. This can be a valid option for other types of noise such as zero-mean, Laplace-distributed noise, identically distributed across all the observations. The remaining covariance matrix to be considered is  $\boldsymbol{\Sigma}_{u,1}$ . In practice, allowing too much flexibility can make EP updates unstable, and although we could force this matrix to be diagonal, forcing it to be isotropic leads to much more stable results. Thus, we define  $\boldsymbol{\Sigma}_{u,1} = c_1 \mathbf{I}$ , where  $\mathbf{I} \in \mathbb{R}^{N \times N}$  denotes an identity matrix. An extended discussion about constraining covariance matrices in EP is included in section 7. Once the constraints on different covariance matrices are set, we can discuss how the problems in (3.6) can be solved. Since (3.6c) is the same problem as (2.6a), we focus on (3.6a) and (3.6b).

*EP update of  $q_{u,0}(\mathbf{u})$ .* In (3.6a), we need to compute the first- and second-order moments of the tilted distribution  $P_0(\mathbf{u}) = \prod_{n=1}^N \bar{\mathcal{P}}_{y_n}(u_n) q_{u,1}(\mathbf{u})$ . Fortunately, this distribution factorizes over the  $N$  elements of  $\mathbf{u}$  since  $q_{u,1}(\mathbf{u})$  has an isotropic (and thus diagonal) covariance matrix. Consequently, we only need to compute the marginal means and variances associated with  $P_0(\mathbf{u})$ , denoted by  $\mathbb{E}_{P_0}[u_n]$  and  $\text{Var}_{P_0}[u_n]$ . These moments can be computed using only 1D integrals by introducing the marginal tilted distributions  $P_0(u_n) = \bar{\mathcal{P}}_{y_n}(u_n) \mathcal{N}(u_n; \mu_{1,n}, c_1)$ , where we use  $\mu_{1,n} := m_{u,1,n}$ ,  $\forall n \in N$  to simplify the notation.

If  $y_n = 0$ , then  $\bar{\mathcal{P}}_{y_n=0}(u_n)$  reduces to  $e^{-u_n} \mathbb{I}(u_n > 0) + \mathbb{I}(u_n \leq 0)$ , and the tilted distribution  $P_0(u_n)$  is a mixture of two truncated Gaussian distributions, whose mean and variance can be computed in closed-form [48, Chap. 2]. On the other hand, if  $y_n \neq 0$ , computing the mean and variance of  $P_0(u_n)$  is less trivial but still possible by using only the following 1D integrals:

$$(3.7) \quad \begin{cases} Z_n = \int P_0(u_n) du_n, \\ \mathbb{E}_{P_0}[u_n] = \frac{1}{Z_n} \int u_n P_0(u_n) du_n, \\ \mathbb{E}_{P_0}[u_n^2] = \frac{1}{Z_n} \int u_n^2 P_0(u_n) du_n. \end{cases}$$

Two main approaches can be adopted to compute these integrals. Ko et al. reported recursive formulas to compute the first two moments of  $P_0(u_n)$  at the cost of  $\mathcal{O}(y_n + 2)$  in [47, 48]. However, these recursive formulas were found to be numerically unstable with moderately high values of  $y_n$ . In this work, we use the quadrature-based implementation proposed in [49, Appendix B], which presents a higher computational complexity but is more stable. Such a quadrature method can also be used for other noise models, as long as the observations in  $\mathbf{y}$  are mutually independent given  $\mathbf{u}$ . In such cases, it is sufficient to find approximate solutions to compute (3.7). For instance, in first-photon or low-photon imaging applications,

the observations can be assumed to follow binomial or geometric distributions [50], and the proposed EP framework could be applied to such problems without major modifications.

Once the marginal means  $\{\mathbb{E}_{P_0}[u_n]\}_{n=1,\dots,N}$  and variances  $\{\text{Var}_{P_0}[u_n]\}_{n=1,\dots,N}$  are computed, the mean and marginal variances of  $q_{u,0}(\mathbf{u})$  can be updated element-wise via

$$(3.8) \quad \frac{1}{c_{0,n}} = \frac{1}{\text{Var}_{P_0}[u_n]} - \frac{1}{c_1}, \quad \mu_{0,n} = c_{0,n} \left( \mathbb{E}_{P_0}[u_n] \left( \frac{1}{c_{0,n}} + \frac{1}{c_1} \right) - \frac{\mu_{1,n}}{c_1} \right),$$

using the fact that  $Q(\mathbf{u})$  has a diagonal covariance matrix (see Appendix A.1). Note that (3.8) does not ensure that  $c_{0,n}$  is positive. In practice, if the value obtained is negative, it is replaced by assigning a large value (e.g.,  $10^8$ ) before computing  $\mu_{0,n}$ .

*EP update of  $q_{x,1}(\mathbf{x})$ .* In (3.6b), the joint tilted distribution is  $P_1(\mathbf{u}, \mathbf{x}) = \prod_{n=1}^N \delta(u_n - \mathbf{h}_n \mathbf{x}) q_{u,0}(\mathbf{u}) q_{x,0}(\mathbf{x})$ . However, since the approximation  $Q(\mathbf{u}, \mathbf{x})$  assumes independence between  $\mathbf{u}$  and  $\mathbf{x}$ , it is sufficient to compute the means and marginal variances of  $P_1(\mathbf{x}) = \int P_1(\mathbf{u}, \mathbf{x}) d\mathbf{u}$  and  $P_1(\mathbf{u}) = \int P_1(\mathbf{u}, \mathbf{x}) d\mathbf{x}$  to update  $q_{u,1}(\mathbf{u})$  and  $q_{x,1}(\mathbf{x})$ , respectively.

To update  $q_{x,1}(\mathbf{x})$ , we need to compute the moments of  $P_1(\mathbf{x}) \propto \int P_1(\mathbf{u}, \mathbf{x}) d\mathbf{u}$ , which can be computed in closed-form using

$$(3.9) \quad \begin{aligned} P_1(\mathbf{x}) &\propto q_{x,0}(\mathbf{x}) \int \delta(\mathbf{u} - \mathbf{H}\mathbf{x}) q_{u,0}(\mathbf{u}) d\mathbf{u} \\ &\propto q_{x,0}(\mathbf{x}) \mathcal{N}(\mathbf{H}\mathbf{x}; \mathbf{m}_{u,0}, \Sigma_{u,0}). \end{aligned}$$

It follows that the mean  $\mathbb{E}_{P_1}[\mathbf{x}]$  and covariance  $\text{Cov}_{P_1}(\mathbf{x})$  are given by

$$(3.10) \quad \text{Cov}_{P_1}(\mathbf{x}) = (\Sigma_{x,0}^{-1} + \mathbf{H}^T \Sigma_{u,0}^{-1} \mathbf{H})^{-1}, \quad \mathbb{E}_{P_1}[\mathbf{x}] = \text{Cov}_{P_1}(\mathbf{x}) \left( \Sigma_{x,0}^{-1} \mathbf{m}_{x,0} + \mathbf{H}^T \Sigma_{u,0}^{-1} \mathbf{m}_{u,0} \right).$$

Since  $Q(\mathbf{x})$  has a block-diagonal covariance matrix, the update of  $\mathbf{m}_{x,1}$  and the blocks on the diagonal of  $\Sigma_{x,1}^{-1}$  is performed the same as in section 2.4, i.e., using the RBMC method.

*EP update of  $q_{u,1}(\mathbf{u})$ .* Since  $Q(\mathbf{u})$  has a diagonal covariance matrix, we only need to compute the marginal moments of  $P_1(u_n) \propto \int P_1(\mathbf{u}) d\mathbf{u}_{\setminus n}, \forall n \in N$ , where  $\mathbf{u}_{\setminus n}$  denotes the elements of  $\mathbf{u}$  whose element  $u_n$  has been removed. Using the formula provided in [33] to compute the integral of the product of a Gaussian density by a Dirac delta function, we obtain

$$(3.11) \quad \begin{aligned} P_1(u_n) &\propto \int \int P_1(\mathbf{u}, \mathbf{x}) d\mathbf{u}_{\setminus n} d\mathbf{x} \\ &\propto \int \int \prod_{n=1}^N \delta(u_n - \mathbf{h}_n \mathbf{x}) q_{u,0}(\mathbf{u}) q_{x,0}(\mathbf{x}) d\mathbf{u}_{\setminus n} d\mathbf{x} \\ &\propto q_{u,0}(u_n) \int \delta(u_n - \mathbf{h}_n \mathbf{x}) \int \delta(\mathbf{u}_{\setminus n} - \mathbf{H}_{\setminus n} \mathbf{x}) q_{u,0}(\mathbf{u}_{\setminus n}) d\mathbf{u}_{\setminus n} q_{x,0}(\mathbf{x}) d\mathbf{x} \\ &\propto q_{u,0}(u_n) \int \delta(u_n - \mathbf{h}_n \mathbf{x}) q_{u,0}^{(\setminus n)}(\mathbf{H}_{\setminus n} \mathbf{x}) q_{x,0}(\mathbf{x}) d\mathbf{x} \\ &\propto q_{u,0}(u_n) \mathcal{N}(u_n; \mathbf{h}_n \tilde{\mathbf{m}}_n, \mathbf{h}_n \tilde{\Sigma}_n^{-1} \mathbf{h}_n^T). \end{aligned}$$

$\mathbf{H}_{\setminus n}$  is the matrix  $\mathbf{H}$  whose  $n$ th row has been removed,  $q_{u,0}^{(\setminus n)}(\mathbf{H}_{\setminus n} \mathbf{x}) = \mathcal{N}(\mathbf{H}_{\setminus n} \mathbf{x}; \mathbf{m}_{u,0}^{\setminus n}, \Sigma_{u,0}^{\setminus n})$ ,  $\tilde{\mathbf{m}}_n = \tilde{\Sigma}_n^{-1} [\mathbf{H}_{\setminus n}^T (\Sigma_{u,0}^{\setminus n})^{-1} \mathbf{m}_{u,0}^{\setminus n} + \Sigma_{x,0}^{-1} \mathbf{m}_{x,0}]$ , and  $\tilde{\Sigma}_n = \mathbf{H}_{\setminus n}^T (\Sigma_{u,0}^{\setminus n})^{-1} \mathbf{H}_{\setminus n} + \Sigma_{x,0}^{-1}$ . The last line

---

**Algorithm 2.** Proposed patch-based EP algorithm—Poisson observation model.

---

**Input:**  $\mathbf{y}$ ,  $\mathbf{H}$ ,  $\{\omega_k, \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k, K\}$

**Output:**  $\boldsymbol{\Sigma}_{u,1}, \mathbf{m}_{u,1}, \boldsymbol{\Sigma}_{u,0}, \mathbf{m}_{u,0}, \boldsymbol{\Sigma}_{x,0}, \mathbf{m}_{x,0}, \boldsymbol{\Sigma}_{x,1}, \mathbf{m}_{x,1}$

```

1: initialization:  $\{\mathbf{m}_{u,0}, \mathbf{m}_{u,1}, \mathbf{m}_{x,1}, \mathbf{m}_{x,0}\} = \mathbf{y} + 1$ ,  $\{\boldsymbol{\Sigma}_{u,0}, \boldsymbol{\Sigma}_{u,1}, \boldsymbol{\Sigma}_{x,1}, \boldsymbol{\Sigma}_{x,0}\} = \text{diag}(\mathbf{y} + 1)$ 
2: while stopping criterion is not satisfied do
3:   • EP update of  $q_{u,0}(\mathbf{u})$ 
4:     compute titled covariance and mean of titled distribution in (3.6a)
5:     update the marginal variance and mean of  $q_{u,0}(\mathbf{u})$  via (3.8)
6:   • EP update of  $q_{x,1}(\mathbf{x})$ 
7:     compute titled mean and block covariance by RBMC in (2.16)
8:     for  $j = 1, \dots, J$  in parallel do
9:       update the precision  $\boldsymbol{\Omega}_{1,j}$  of diagonal blocks
10:      update the covariance of diagonal block  $[\boldsymbol{\Sigma}_{x,1}]_j$  and subset of mean  $[\mathbf{m}_{x,1}]_j$ 
11:     end for
12:   • EP update of  $q_{u,1}(\mathbf{u})$ 
13:     compute the marginal moments of  $\{P_1(u_n)\}_{n=1, \dots, N}$  in (3.12)
14:     update the mean and variance of  $q_{u,1}(\mathbf{u})$  via (3.13)
15:   • EP update of  $q_{x,0}(\mathbf{x})$ 
16:     for  $j = 1, \dots, J$  in parallel do
17:       update the precision  $\boldsymbol{\Omega}_{0,j}$  of diagonal blocks
18:       update the covariance of diagonal block  $[\boldsymbol{\Sigma}_{x,0}]_j$  and subset of mean  $[\mathbf{m}_{x,0}]_j$ 
19:     end for
20: end while

```

---

of (3.11) shows that  $P_1(u_n)$  is the product of two Gaussian densities. Thus its moments are given by

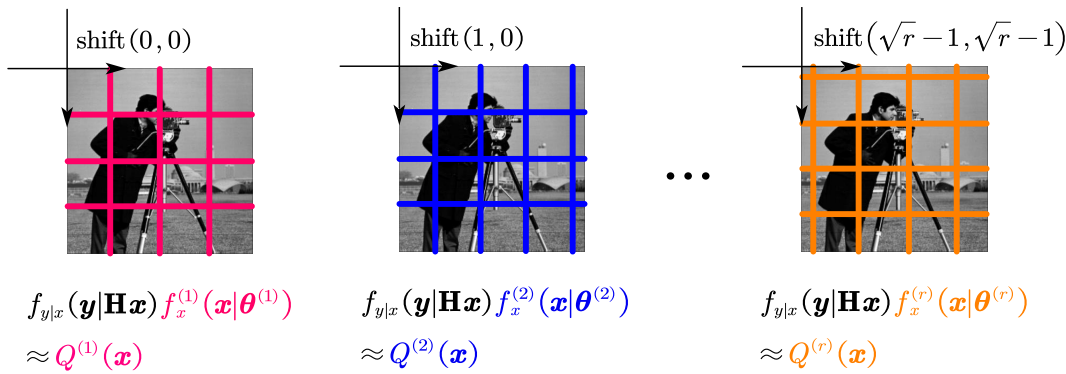
$$(3.12) \quad \text{Var}_{P_1}(u_n) = \left( \frac{1}{c_{0,n}} + \frac{1}{\mathbf{h}_n \tilde{\boldsymbol{\Omega}}_n^{-1} \mathbf{h}_n^T} \right)^{-1}, \quad \mathbb{E}_{P_1}[u_n] = \text{Var}_{P_1}(u_n) \left( \frac{\mu_{0,n}}{c_{0,n}} + \frac{\mathbf{h}_n \tilde{\mathbf{m}}_n}{\mathbf{h}_n \tilde{\boldsymbol{\Omega}}_n^{-1} \mathbf{h}_n^T} \right).$$

The main computational cost here is the inversion of the large pixel-dependent matrix  $\tilde{\boldsymbol{\Omega}}_n^{-1}$  needed to compute  $\mathbf{h}_n \tilde{\mathbf{m}}_n$  and  $\mathbf{h}_n \tilde{\boldsymbol{\Omega}}_n^{-1} \mathbf{h}_n^T$ . However, it can be noticed that  $\tilde{\boldsymbol{\Omega}}_n^{-1}$  is  $\text{Cov}_{P_1}(\mathbf{x})$  in (3.10) whose contribution of  $u_n$  has been removed. Even for relatively small images,  $\tilde{\boldsymbol{\Omega}}_n^{-1} \approx \text{Cov}_{P_1}(\mathbf{x})$  is a reasonable approximation. Instead of using this approximation which would still require the computation of  $\text{Cov}_{P_1}(\mathbf{x})$ , we propose using  $\tilde{\boldsymbol{\Omega}}_n^{-1} \approx \boldsymbol{\Sigma}_*$ , which is a block-diagonal approximation of  $\text{Cov}_{P_1}(\mathbf{x})$  (computed during the update of  $q_{x,1}(\mathbf{x})$  above). This allows a fast approximation of the marginal moments in (3.12). Using these moments and  $q_{u,0}(\mathbf{u})$ , the updated  $c_1$  can be obtained via gradient descent (or Newton–Raphson), as detailed in Appendix A.2. After the update of  $c_1$ ,  $\mathbf{m}_{u,1}$  is obtained via

$$(3.13) \quad \mathbf{m}_{u,1} = c_1 \left[ \left( c_1^{-1} \mathbf{I} + \boldsymbol{\Sigma}_{u,0}^{-1} \right) \mathbb{E}_{P_1}[\mathbf{u}] - \boldsymbol{\Sigma}_{u,0}^{-1} \mathbf{m}_{u,0} \right].$$

The pseudo-code of the EP algorithm for the Poisson observation model is presented in Algorithm 2. The updates of  $q_{u,0}(\mathbf{u})$ ,  $q_{x,1}(\mathbf{x})$ ,  $q_{u,1}(\mathbf{u})$ , and  $q_{x,0}(\mathbf{x})$  are repeated until some





**Figure 3.** Illustration of a set of different GMM patch-based priors  $f_x^{(i)}(\mathbf{x}|\boldsymbol{\theta}^{(i)})$  ( $i = 1, \dots, r$ ) modeling overlapping patches extracted from the same image. Approximation for exact posterior  $f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$  is finally built as the Product-of-Experts  $Q_f(\mathbf{x}) \propto [Q^{(1)}(\mathbf{x})Q^{(2)}(\mathbf{x}) \dots Q^{(r)}(\mathbf{x})]^{\frac{1}{r}}$ .

convergence conditions are met. Here, we use  $\|\mathbf{m}_*^{(t)} - \mathbf{m}_*^{(t-1)}\|^2 < 1e^{-8}N$  and  $\|\text{diag}(\boldsymbol{\Sigma}_*^{(t)}) - \text{diag}(\boldsymbol{\Sigma}_*^{(t-1)})\|^2 < 1e^{-8}N$  as the stopping criteria ( $t$  denotes the  $t$ th EP iteration). In general, there are no convergence guarantees for EP algorithms. A damping strategy [27, 31, 51] is often used when oscillations between successive iterations are observed. Here, a damping with parameter  $\epsilon = 0.7$  is applied in Algorithms 1 and 2. After the last iteration, irrespective of the actual noise model, the EP approximation of the posterior  $f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta})$  is given by  $Q(\mathbf{x}) \propto q_{x,0}(\mathbf{x})q_{x,1}(\mathbf{x})$ , with mean  $\mathbf{m}_*$  and covariance matrix  $\boldsymbol{\Sigma}_*$ .

**4. Final estimation using Product-of-Experts.** Using the GMM prior in (2.2), the mean of  $Q(\mathbf{x})$  presents blocky artifacts because the patches are non-overlapping and a priori mutually independent. To tackle this problem, instead of using a single prior and a unique partitioning of the pixels, we consider a set of prior models, with shifted partitions, as illustrated in Figure 3. Using vertical and horizontal steps of one pixel, we build  $r$  distinct partitions, denoted  $\mathcal{S}_1, \dots, \mathcal{S}_r$ , leading to  $r$  prior models  $\{f_x^{(i)}(\mathbf{x}|\boldsymbol{\theta}^{(i)})\}_{i=1, \dots, r}$ . For each of these models, the prior covariance matrix presents (after appropriate sorting of the pixels in  $\mathbf{x}$ ) a block-diagonal structure. Thus, the EP methods described in sections 2 and 3 can be used separately with any of these priors, which provide a set of  $r$  approximations  $Q^{(i)}(\mathbf{x}) \approx f(\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}^{(i)})$ . These distributions can be combined in different ways. Model averaging could be achieved by assigning a prior probability (denoted  $f(\mathcal{S}_i)$ , e.g.,  $f(\mathcal{S}_i) = 1/r$ ) to each partitioning of the pixels, and by marginalizing over the  $r$  partitions [52, Chapter 7.4]. This would require approximating the evidences  $f(\mathbf{y}|\boldsymbol{\theta}^{(i)})$  which is possible within EP [25]. However, this approach does not, in general, provide reconstructed images of satisfactory quality as only a few partitions effectively contribute to the final result and blocky artifacts remain present. Similarly, the simple Gaussian mixture  $\sum_{i=1}^r \frac{1}{r} Q^{(i)}(\mathbf{x})$  also leads to poor final MMSE estimates and large posterior variances. Here, we adopt an alternative strategy which does not require the computation/approximation of model evidences, whereby each approximation  $Q^{(i)}(\mathbf{x})$  is seen as an expert, and the final distribution  $Q_f(\mathbf{x})$  representing  $\mathbf{x}$  is built as the Product-of-Experts

(PoE) [53], i.e.,

$$(4.1) \quad Q_f(\mathbf{x}) \propto \left[ \prod_{i=1}^r Q^{(i)}(\mathbf{x}) \right]^{\frac{1}{r}}.$$

Each approximation  $Q^{(i)}(\mathbf{x})$  is a multivariate Gaussian distribution with mean and covariance matrix denoted by  $\mathbf{m}_*^{(i)}$  and  $\Sigma_*^{(i)}$ , and so is  $Q_f(\mathbf{x})$ . Its mean and covariance matrix are given by

$$(4.2) \quad \text{Cov}_{Q_f(\mathbf{x})}(\mathbf{x}) = \left[ \frac{1}{r} \sum_{i=1}^r (\Sigma_*^{(i)})^{-1} \right]^{-1}, \quad \mathbb{E}_{Q_f(\mathbf{x})}[\mathbf{x}] = \text{Cov}_{Q_f(\mathbf{x})}(\mathbf{x}) \left[ \frac{1}{r} \sum_{i=1}^r (\Sigma_*^{(i)})^{-1} \mathbf{m}_*^{(i)} \right].$$

While it is possible to use the individual block-diagonal covariance matrices  $\Sigma_*^{(i)}$  in (4.2), the resulting marginal variances of  $Q_f(\mathbf{x})$  are often smaller than expected. Instead, we only use the marginal variances of the experts to compute the marginal moments of  $Q_f(\mathbf{x})$ , and the final image point estimate considered is  $\mathbb{E}_{Q_f(\mathbf{x})}[\mathbf{x}]$ , the mean of  $Q_f(\mathbf{x})$ . Although the PoE strategy adopted in this paper differs from more traditional image priors based on overlapping patches, the results presented in section 6 show that the resulting estimated images are on par with existing patch-based methods and that the estimated posterior variances, despite being approximated, can provide useful information about the posterior uncertainty measures of the image intensities.

**5. EP-EM strategy for hyperparameter estimation.** So far, we have treated  $\theta$  in the prior model  $f_x(\mathbf{x}|\theta) = \prod_{j=1}^J \sum_{k=1}^K \omega_k \mathcal{N}(\mathbf{x}_j; \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k)$  as a set of known hyperparameters. In this section, with a fixed  $K$  in  $\theta = \{\omega_k, \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k, K\}$ , a subset of the parameters in  $\theta$  is pre-trained using external images, while another subset of the parameters is retrieved in an unsupervised manner via an Expectation Maximization (EM) scheme leveraging the Gaussian approximations of the proposed EP algorithms. This is motivated by the fact that training patch-based GMMs in an unsupervised fashion, based on data potentially corrupted by non-Gaussian noise and by a linear degradation operator  $\mathbf{H}$  (especially convolution operators), is extremely difficult. On the other hand, it is easy to train patch-based GMMs using many external normalized patches and then only estimate a reduced set of scaling/normalization parameters for an individual image. Without loss of generality, a conventional GMM can be trained using an EM algorithm on a set of external clean image patches [5] (or even noisy images [8, 9]). However, such images/patches are often normalized beforehand [5, 8, 9, 10]; e.g., the pixel intensities are scaled to  $[0, 1]$ , and the means of the patches are removed prior to training the GMM. Let  $\tilde{\mathbf{x}}_j$  represent a zero-mean image patch extracted from an image whose intensity is scaled in  $[0, 1]$ . A conventional GMM can be expressed as

$$(5.1) \quad f(\tilde{\mathbf{x}}_j | \{\omega_k, \boldsymbol{\mu}_k, \mathbf{C}_k\}_{k=1, \dots, K}) = \sum_{k=1}^K \omega_k \mathcal{N}(\tilde{\mathbf{x}}_j; \boldsymbol{\mu}_k, \mathbf{C}_k)$$

and can be characterized by  $\{\omega_k, \boldsymbol{\mu}_k, \mathbf{C}_k\}_{k=1, \dots, K}$ . Such a GMM can be trained using a set of patches  $\{\tilde{\mathbf{x}}_j\}_{j=1, \dots, J}$ . When denoising/inpainting arbitrary patches using  $\{\omega_k, \boldsymbol{\mu}_k, \mathbf{C}_k\}_{k=1, \dots, K}$ ,

the mean of each noisy patch can be removed first, the resulting centered patches are then restored, and finally the original patches' means are added back. In more complex scenarios, this empirical approach has limited interest as adding/removing patch means can complicate the resulting noise statistics. If the GMM is trained without patch mean subtraction, a much larger number of components should be used. Moreover, regardless of whether the patch means are subtracted or not, the trained GMM (in particular its covariance matrices) depends significantly on the scale of the images used for training. While in the Gaussian noise case, it is possible to rescale the noisy observations so that the intensities match those of the training images, this is not possible in the Poisson case as the resulting observation noise would no longer be Poisson distributed. To solve these two problems, we propose to adjust the GMM prior by incorporating and estimating an offset and a scale parameter. More precisely, we express patches of  $\mathbf{x}$  as  $\mathbf{x}_j = \bar{x}_j \mathbf{1} + \alpha \tilde{\mathbf{x}}_j$ , where  $\bar{x}_j$  denotes the mean of the patch  $\mathbf{x}_j$ ,  $\tilde{\mathbf{x}}_j$  is a zero-mean patch which is assumed to follow the GMM in (5.1), and  $\alpha > 0$  is a scaling factor. It follows that

$$(5.2) \quad f(\mathbf{x}_j | \bar{x}_j, \alpha, \{\omega_k, \boldsymbol{\mu}_k, \mathbf{C}_k\}_{k=1, \dots, K}) = \sum_{k=1}^K \omega_k \mathcal{N}(\mathbf{x}_j; \bar{x}_j \mathbf{1} + \alpha \boldsymbol{\mu}_k, \alpha^2 \mathbf{C}_k).$$

The mean  $\bar{x}_j$  is unknown in practice. It can be approximated by the mean of the noisy patches, but this is prone to errors in particular if the observed image is blurred, and/or corrupted by high levels of noise. Instead of fixing  $\bar{x}_j$ , it is assigned a Gaussian prior distribution and shared across all the patches; i.e.,  $f(\bar{x}_j) = \mathcal{N}(\bar{x}_j; m_0, s^2)$ ,  $\forall j = 1, \dots, J$ . This simple prior model allows  $\bar{x}_j$  to be marginalized, leading to the final GMM prior

$$(5.3) \quad \begin{aligned} f_x(\mathbf{x}_j | m_0, s^2, \alpha, \{\omega_k, \boldsymbol{\mu}_k, \mathbf{C}_k\}_{k=1, \dots, K}) &= \sum_{k=1}^K \omega_k \mathcal{N}(\mathbf{x}_j; m_0 \mathbf{1} + \alpha \boldsymbol{\mu}_k, s^2 \mathbf{1} \mathbf{1}^T + \alpha^2 \mathbf{C}_k) \\ &= \sum_{k=1}^K \omega_k \mathcal{N}(\mathbf{x}_j; \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k) = f_x(\mathbf{x}_j | \boldsymbol{\theta}), \end{aligned}$$

where  $\tilde{\boldsymbol{\mu}}_k = m_0 \mathbf{1} + \alpha \boldsymbol{\mu}_k$  and  $\tilde{\mathbf{C}}_k = s^2 \mathbf{1} \mathbf{1}^T + \alpha^2 \mathbf{C}_k$ . In this work, the parameters  $\{\omega_k, \boldsymbol{\mu}_k, \mathbf{C}_k, K\}$  of the trained GMM were taken from [5], and these parameters were fixed, while the remaining hyperparameters  $\boldsymbol{\theta}^\dagger = \{m_0, s^2, \alpha\}$  are estimated.

A classical approach to hyperparameter estimation is via Expectation Maximization (EM) [54], whereby  $\mathbf{x}$  is treated as a latent vector and  $\boldsymbol{\theta}^\dagger$  is estimated by maximizing the marginal likelihood  $f(\mathbf{y} | \boldsymbol{\theta}^\dagger)$  or marginal posterior distribution  $f(\boldsymbol{\theta}^\dagger | \mathbf{y})$ . Unfortunately, a simple EM approach is not possible here since it would require computing expectation w.r.t. the exact posterior  $f(\mathbf{x} | \mathbf{y}, \boldsymbol{\theta}^\dagger)$ . However, it is possible to use a variational EM method [55], whereby the expectations w.r.t.  $f(\mathbf{x} | \mathbf{y}, \boldsymbol{\theta}^\dagger)$  are replaced by expectations w.r.t. an alternative approximating distribution (i.e., the EP approximation here), leading to an EP-EM scheme similar to that used in [56]. Since  $\boldsymbol{\theta}^\dagger = \{m_0, s^2, \alpha\}$  contains only three parameters shared across all the  $J$  patches, we do not assign  $\boldsymbol{\theta}^\dagger$  a hyperprior, as in practice  $f(\mathbf{y} | \boldsymbol{\theta}^\dagger)$  seen as a function of  $\boldsymbol{\theta}^\dagger$  is sufficiently concentrated around its mode.

As mentioned in section 4, one of the main reasons for considering independent experts is that the computational load can be distributed. Although  $\boldsymbol{\theta}^\dagger$  could be shared across all the

experts, here each expert has its own set of hyperparameters  $\boldsymbol{\theta}^{(\dagger,i)}$  ( $i = 1, \dots, r$ ) which are estimated independently. To simplify the notation, in the remainder of this section, the expert indices are omitted. The natural joint likelihood to be considered within EP-EM is  $f(\mathbf{x}, \mathbf{y} | \boldsymbol{\theta}^\dagger)$ , but this requires the computation of the expected value of the log-prior  $\log(f(\mathbf{x} | \boldsymbol{\theta}^\dagger))$ , i.e., the logarithm of a weighted sum of multivariate Gaussian p.d.f.s, which makes the estimation of  $\boldsymbol{\theta}^\dagger$  intractable. A more suitable GMM formulation consists of introducing, for each patch, a discrete auxiliary variable  $z_j \in \{1, \dots, K\}$  [23] and the following hierarchical prior model:

$$(5.4) \quad f(\mathbf{x} | \mathbf{z}, \boldsymbol{\theta}^\dagger) = \prod_{j=1}^J \prod_{k=1}^K \left[ \mathcal{N}(\mathbf{x}_j; \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k) \right]^{\delta(z_j-k)}, \quad f(\mathbf{z}) = \prod_{j=1}^J f(z_j),$$

with  $\mathbf{z} = [z_1, \dots, z_J]^T$  and  $f(z_j = k) = \omega_k$ ,  $\forall j \in \{1, \dots, J\}$ . By marginalizing  $\mathbf{z}$  in (5.4), we recover the model (2.2). To simplify the description of the EP methods in sections 2 and 3, we discussed approximations of the exact posterior  $f(\mathbf{x} | \mathbf{y}, \boldsymbol{\theta}^\dagger)$  (section 2) and  $f(\mathbf{x}, \mathbf{u} | \mathbf{y}, \boldsymbol{\theta}^\dagger)$  (section 3). However, EP can also be used to approximate the extended posterior distributions

$$(5.5) \quad \begin{aligned} f(\mathbf{x}, \mathbf{z} | \mathbf{y}, \boldsymbol{\theta}^\dagger) &\propto f(\mathbf{y} | \mathbf{x}) f(\mathbf{x} | \mathbf{z}, \boldsymbol{\theta}^\dagger) f(\mathbf{z}) && \text{(Gaussian noise case),} \\ f(\mathbf{x}, \mathbf{u}, \mathbf{z} | \mathbf{y}, \boldsymbol{\theta}^\dagger) &\propto f(\mathbf{y} | \mathbf{u}) f(\mathbf{u} | \mathbf{x}) f(\mathbf{x} | \mathbf{z}, \boldsymbol{\theta}^\dagger) f(\mathbf{z}) && \text{(Poisson noise case).} \end{aligned}$$

This can be achieved easily using the EP factorization used in [31], and it leads to an EP approximation that factorizes as  $Q^\dagger(\mathbf{x})Q^\dagger(\mathbf{z})$  (resp.,  $Q^\dagger(\mathbf{x})Q^\dagger(\mathbf{u})Q^\dagger(\mathbf{z})$ ) in the Gaussian noise case (resp., Poisson noise case). Note that we use the notation  $Q^\dagger(\cdot)$  to highlight the approximations related to the extended models in (5.5). Irrespective of the noise model considered,  $Q^\dagger(\mathbf{z})$  is chosen to be a product of  $J$  independent categorical distributions, one for each element of  $\mathbf{z}$ , which can take  $K$  values. It turns out that, for the Gaussian case,  $Q^\dagger(\mathbf{x})$  and  $Q^\dagger(\mathbf{z})$  in the product above can be obtained from our EP method which approximates  $f(\mathbf{x} | \mathbf{y}, \boldsymbol{\theta}^\dagger)$ . Indeed, it can be shown that  $Q^\dagger(\mathbf{x}) = Q(\mathbf{x})$  and that the probabilities involved in  $Q^\dagger(\mathbf{z})$  are the weights computed in (2.9), i.e.,  $Q^\dagger(z_j = k) = \hat{\omega}_{j,k}$ . Similarly, in the Poisson noise case, we have  $Q^\dagger(\mathbf{x})Q^\dagger(\mathbf{u}) = Q(\mathbf{x})Q(\mathbf{u})$ . This can be explained by the fact that EP aims to approximate marginal posteriors in both cases (i.e.,  $Q(\mathbf{x}) \approx \int f(\mathbf{x}, \mathbf{u} | \mathbf{y}, \boldsymbol{\theta}^\dagger) d\mathbf{u} = f(\mathbf{x} | \mathbf{y}, \boldsymbol{\theta}^\dagger)$  and  $Q(\mathbf{x})^\dagger \approx \sum_{\mathbf{z}} \int f(\mathbf{x}, \mathbf{u}, \mathbf{z} | \mathbf{y}, \boldsymbol{\theta}^\dagger) d\mathbf{u} = f(\mathbf{x} | \mathbf{y}, \boldsymbol{\theta}^\dagger)$ ), and since both EP schemes rely on the same factorization, with the same covariance constraints, the marginal approximations are the same.

Using the extended likelihood  $f(\mathbf{y}, \mathbf{x}, \mathbf{z} | \boldsymbol{\theta}^\dagger)$  (or  $f(\mathbf{y}, \mathbf{x}, \mathbf{u}, \mathbf{z} | \boldsymbol{\theta}^\dagger)$  in the Poisson case), the resulting EP-EM reduces to two sequential steps. As mentioned in section 2,  $Q(\mathbf{x})$  implicitly depends on  $\boldsymbol{\theta}^\dagger$  (and  $\mathbf{y}$ ). While the dependence on  $\boldsymbol{\theta}^\dagger$  was not explicit so far to simplify the notation, we now use  $Q^\dagger(\cdot | \boldsymbol{\theta}^\dagger)$  as  $\boldsymbol{\theta}^\dagger$  is no longer fixed. The ( $t$ )th iteration of EM includes the following *E-step* and *M-step*.

**E-step.** When computing the expected value of the extended likelihood, the only factor that depends on  $\boldsymbol{\theta}^\dagger$  is  $f(\mathbf{x} | \mathbf{z}, \boldsymbol{\theta}^\dagger)$ , leading to the cost function

$$C(\boldsymbol{\theta}^\dagger | \boldsymbol{\theta}_{(t-1)}^\dagger) = \mathbb{E}_{Q^\dagger(\mathbf{x}, \mathbf{z} | \boldsymbol{\theta}_{(t-1)}^\dagger)} \left[ \log f(\mathbf{x} | \mathbf{z}, \boldsymbol{\theta}^\dagger) \right] = \sum_{j,k} \hat{\omega}_{j,k} \mathbb{E}_{Q^\dagger(\mathbf{x} | \boldsymbol{\theta}_{(t-1)}^\dagger)} \left[ \log \mathcal{N}(\mathbf{x}_j; \tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k) \right],$$

**Algorithm 3.** Full description for the proposed patch-based EP algorithm.

**Input:**  $\mathbf{y}$ ,  $\mathbf{H}$ ,  $\sigma^2$  (if Gaussian i.i.d. noise), conventional trained GMM  $\{\omega_k, \boldsymbol{\mu}_k, \mathbf{C}_k, K\}$

**Output:**  $\mathbb{E}_{Q_f(\mathbf{x})}[\mathbf{x}]$ ,  $\text{Cov}_{Q_f(\mathbf{x})}(\mathbf{x})$

- 1: **for**  $i = 1: r$  in parallel **do**
- 2:     **while** stopping criterion is not satisfied **do**
- 3:         • Compute  $Q^{(\dagger,i)}(\mathbf{x}|\boldsymbol{\theta}^{\dagger,i})$  using Algorithm 1 (for Gaussian observation model) or
- 4:         Algorithm 2 (for Poisson observation model)
- 5:         • Estimate  $\boldsymbol{\theta}^{\dagger,i} = \{m_0, s^2, \alpha\}^{\dagger,i}$
- 6:     **end while**
- 7:     **return:**  $\mathbf{m}_*^{(i)}$ ,  $\boldsymbol{\Sigma}_*^{(i)}$ ,  $\boldsymbol{\theta}^{\dagger,i}$
- 8: **end for**
- 9: Compute final moments via (4.2) to obtain  $\text{Cov}_{Q_f(\mathbf{x})}(\mathbf{x})$ ,  $\mathbb{E}_{Q_f(\mathbf{x})}[\mathbf{x}]$

where  $\boldsymbol{\theta}_{(t-1)}^\dagger$  is the value of  $\boldsymbol{\theta}^\dagger$  estimated at the  $(t-1)$ th iteration,  $\hat{\omega}_{j,k}$  is estimated in (2.9), and  $(\tilde{\boldsymbol{\mu}}_k, \tilde{\mathbf{C}}_k)$  are defined as in (5.3). Since  $Q^\dagger(\mathbf{x}|\boldsymbol{\theta}_{(t-1)}^\dagger)$  is a Gaussian distribution with mean  $\mathbf{m}_*$  and covariance matrix  $\boldsymbol{\Sigma}_*$ , the cost function takes the form

$$C(\boldsymbol{\theta}^\dagger|\boldsymbol{\theta}_{(t-1)}^\dagger) = -\frac{1}{2} \sum_{k,j} \hat{\omega}_{j,k} \log |s^2 \mathbf{1}\mathbf{1}^T + \alpha^2 \mathbf{C}_k| - \frac{1}{2} \sum_{j,k} \hat{\omega}_{j,k} \text{tr}[(s^2 \mathbf{1}\mathbf{1}^T + \alpha^2 \mathbf{C}_k)^{-1} [\boldsymbol{\Sigma}_*]_j] + c$$

$$-\frac{1}{2} \sum_{j,k} \hat{\omega}_{j,k} [([\mathbf{m}_*]_j - m_0 \mathbf{1} - \alpha \boldsymbol{\mu}_k)^T (s^2 \mathbf{1}\mathbf{1}^T + \alpha^2 \mathbf{C}_k)^{-1} ([\mathbf{m}_*]_j - m_0 \mathbf{1} - \alpha \boldsymbol{\mu}_k)],$$

where  $c = -(Jr/2) \log(2\pi)$ .

**M-step.** Estimate  $\boldsymbol{\theta}_{(t)}^\dagger$  by solving  $\boldsymbol{\theta}_{(t)}^\dagger = \text{argmax}_{\boldsymbol{\theta}^\dagger} C(\boldsymbol{\theta}^\dagger|\boldsymbol{\theta}_{(t-1)}^\dagger)$ . For a given value of  $(s^2, \alpha)$ , the optimal value of  $m_0$  can be obtained analytically by

$$(5.6) \quad \hat{m}_0(s^2, \alpha) = \frac{\sum_{j,k} \hat{\omega}_{j,k} ([\mathbf{m}_*]_j - \alpha \boldsymbol{\mu}_k)^T (s^2 \mathbf{1}\mathbf{1}^T + \alpha^2 \mathbf{C}_k)^{-1} \mathbf{1}}{\sum_{j,k} \hat{\omega}_{j,k} \mathbf{1}^T (s^2 \mathbf{1}\mathbf{1}^T + \alpha^2 \mathbf{C}_k)^{-1} \mathbf{1}}.$$

Consequently, it is sufficient to optimize  $(s^2, \alpha)$ . In practice, instead of an exhaustive search on a 2D grid for  $(s^2, \alpha)$ , we first fix  $\alpha$  and update  $s^2$  and then fix  $s^2$  and optimize  $\alpha$ . This procedure is repeated until convergence. The pseudo-code of the final EP-EM algorithm is summarized in Algorithm 3. It is important to mention that estimating only the three parameters  $(m_0, s^2, \alpha)$  instead of all the weights, means, and covariance matrices of the  $K$ -component GMM allows for a more reliable hyperparameter estimation. Since these three parameters are shared by the  $K$  components of the GMM, their estimation is usually reliable, irrespective of the number of components in the mixture.

Note that in practice, the image to be reconstructed may not be square, or its dimension may not match the size of the GMM patch-based prior. It is however possible to apply the same GMM patch-based prior without additional training, by simply marginalizing elements (e.g., rows or columns) in (5.3) to create GMMs for sub-patches.

**6. Experiments.** In this section, we demonstrate the performance of the proposed patch-based EP algorithm on different image restoration tasks: image denoising, mild image inpainting, and supervised image deconvolution, assuming in each case that the degradation operator  $\mathbf{H}$  is known. The experiments are conducted on synthetic grayscale images for which the ground truth image  $\mathbf{x} \in \mathbb{R}^N$  is known. Two types of operator  $\mathbf{H} \in \mathbb{R}^{N \times N}$  are generated.

- $\mathbf{H}$  is diagonal: it is either (i) an identity matrix (for denoising) or (ii) a diagonal mask matrix with diagonal elements equal to 1 if pixels are observed and 0 otherwise (for inpainting). These operators do not induce correlation between pixels, and the covariance matrices of the approximating factors  $q_{x,1}(\mathbf{x})$ ,  $q_{x,0}(\mathbf{x})$  are kept diagonal. Although spatial correlations are induced by the patch-based prior, maintaining  $q_{x,1}(\mathbf{x})$  and  $q_{x,0}(\mathbf{x})$  as diagonal helps to reduce the computational cost in (2.9), and it is found in practice that the performance is not significantly affected by imposing such constraints for denoising and inpainting. We also demonstrate the benefits of the hyperparameter estimation in these two tasks.
- $\mathbf{H}$  is a 2D convolution matrix, which induces spatial correlation across nearby pixels. In this case, the covariance matrices of the approximating factors  $q_{x,1}(\mathbf{x})$ ,  $q_{x,0}(\mathbf{x})$  are block-diagonal.

As mentioned in section 5, the GMM used here to build the prior models is taken from [5] (available at <https://people.csail.mit.edu/danielzoran/>). It was trained using a set of  $2 \times 10^6$  patches from a generic external clean dataset [57], and it consists of  $K = 200$  components' patches of  $8 \times 8$  pixels. The images restored using EP-EM are obtained via the mean of  $Q_f(\mathbf{x})$ , i.e.,  $\hat{\mathbf{x}}_{\text{proposed}} = \mathbb{E}_{Q_f(\mathbf{x})}[\mathbf{x}]$ . To evaluate quantitatively the quality of a restored image  $\hat{\mathbf{x}}$ , we use the peak signal-to-noise ratio (PSNR) defined as

$$(6.1) \quad \text{PSNR}(\mathbf{x}, \hat{\mathbf{x}}) = 10 \times \log_{10} \frac{(\max(\mathbf{x}))^2}{\frac{1}{N} \|\hat{\mathbf{x}} - \mathbf{x}\|^2}.$$

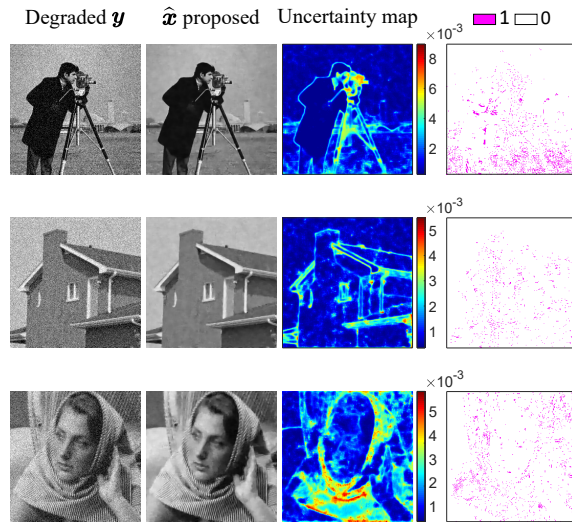
To visualize the pixel-wise approximate uncertainty associated with  $\hat{\mathbf{x}}_{\text{proposed}}$ , we consider the marginal variances of  $\text{Cov}_{Q_f(\mathbf{x})}(\mathbf{x})$ . Apart from the denoising/inpainting problems with Gaussian/Poisson noise for which the EP approximation coincides with the actual posterior distribution, it is difficult to assess the reliability of the approximate posterior uncertainties, as the computation of the exact posterior uncertainties is not tractable. To investigate how far actual pixel values are from the estimated region of high-posterior density, we compute the pixel-wise intervals covering 95% (2.5%–97.5%) of the approximate marginal posteriors, centered around the estimated mean. Binary images indicating the pixels for which the ground truth intensity falls into such an interval are then produced, and we also report the fraction of pixels in that binary image, for which the ground truth intensities lie in those intervals.

All the experiments have been implemented in MATLAB on an Intel(R) Core(TM) i7-8700K CPU @ 3.70GHz workstation. The proposed EP algorithm was performed without parallel implementation, i.e., the  $J$  image patches and  $r$  experts are processed using a sequential *for loop*, and execution times are reported for comparison with the competing methods of each scenario.

**6.1. Gaussian i.i.d. noise model.** This subsection demonstrates the effectiveness of the proposed patch-based EP algorithm when the observation noise is Gaussian and i.i.d. with

	LIDIA	BM3D	MAP-EPLL	Afonso et al.	Proposed
Cameraman (256 × 256)					
10/255	–	<b>34.12</b>	33.99	33.94	<u>34.04</u>
15/255	<b>32.41</b>	<u>31.90</u>	31.79	31.65	31.71
20/255	–	<b>30.45</b>	<u>30.36</u>	30.10	30.19
25/255	<b>29.91</b>	<u>29.21</u>	29.04	28.77	28.83
30/255	–	<b>28.61</b>	<u>28.34</u>	27.99	28.09
50/255	<b>26.83</b>	<u>25.39</u>	25.08	24.55	24.52
House (256 × 256)					
10/255	–	<b>36.79</b>	35.77	35.79	<u>35.82</u>
15/255	<b>35.09</b>	<u>34.97</u>	34.18	34.06	34.12
20/255	–	<b>33.83</b>	<u>33.05</u>	32.75	32.82
25/255	<b>33.08</b>	<u>32.91</u>	32.14	31.66	31.72
30/255	–	<b>32.08</b>	<u>31.25</u>	30.60	30.68
50/255	<b>30.14</b>	<u>29.45</u>	28.91	27.91	27.82
Barbara (256 × 256)					
10/255	–	<b>34.63</b>	33.70	33.79	<u>33.82</u>
15/255	<b>32.90</b>	<u>32.64</u>	31.47	31.41	31.42
20/255	–	<b>31.22</b>	<u>29.91</u>	29.68	29.66
25/255	<u>30.02</u>	<b>30.05</b>	28.63	28.29	28.23
30/255	–	<b>29.09</b>	<u>27.63</u>	27.25	27.18
50/255	<b>26.08</b>	<u>25.97</u>	24.51	24.17	24.02

(a) PSNR (dB) comparison.



(b) Images denoised by the proposed EP-EM algorithm ( $\sigma = 25/255$ ).

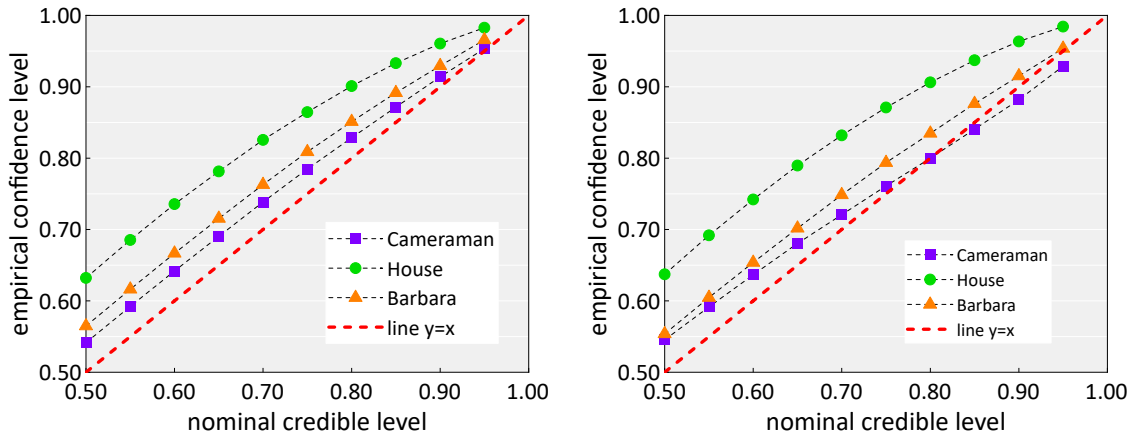
**Figure 4.** Denoising results for Gaussian i.i.d. noise. (a) PSNR (dB) comparison of different denoising methods. The highest and second highest PSNR values are in bold and underlined, respectively. (b) Restored images obtained by the proposed EP-EM algorithm (second column), associated uncertainty maps (third column), and binary maps indicating pixels for which ground truth intensity falls into a 2.5%–97.5% credible interval (last column). Pixels are depicted in white (“0”) if the intensity falls into the credible region and in pink (“1”) otherwise.

known variance  $\sigma^2$ .

**6.1.1. Image denoising and mild inpainting.** Three benchmark grayscale images (*Cameraman*, *House*, *Barbara*) of size  $256 \times 256$  pixels are used as illustrative examples for image denoising and mild inpainting. The pixel intensities are scaled between 0 and 1 so that the same conventional GMM  $\{\omega_k, \mu_k, \mathbf{C}_k, K\}$  can be used within all the competing patch-based methods [5, 8].

For the denoising experiments, the noise standard deviation is chosen from the set  $\sigma \in \{10, 15, 20, 25, 30, 50\}/255$ , which corresponds to PSNR values about  $\{28, 25, 22, 20, 19, 15\}$ dB for the three images. The denoising performance of the proposed EP algorithm is compared to several recent denoising methods assuming Gaussian noise. LIDIA [17] is taken as a benchmark to see how the proposed method differs from state-of-the-art deep learning methods, while more conventional popular patch-based methods are also considered, i.e., BM3D [58], the GMM patch-based method proposed in [8] and referred to as “Afonso et al.,” and the GMM patch-based method using MAP estimation proposed in [5], referred to as “MAP-EPLL”.

The table in Figure 4a compares the PSNR values obtained by different methods. It can be observed that the proposed EP algorithm obtains PSNRs slightly worse than BM3D and LIDIA. Using the same trained GMM as patch-based prior, the proposed EP algorithm provides PSNRs close to those of the MMSE patch-based method [8] and the MAP patch-based method [5], illustrating its effectiveness as an approximate MMSE denoiser. Recall



(a) Denoising with Gaussian i.i.d. noise ( $\sigma = 25/255$ ). (b) Denoising with Poisson noise (peak value = 30).

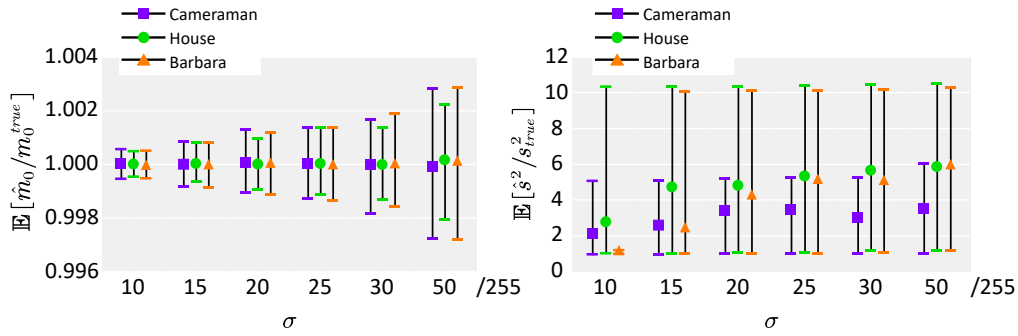
**Figure 5.** Fraction of pixels inside the different credible intervals for image restored by the proposed EP algorithm.

that the main goal of the proposed EP algorithm is not necessarily to achieve the highest PSNRs (although high PSNRs are preferred), but instead to propose a scalable EP framework with patch-based prior. Still, it can provide comparable PSNRs for denoised images with additional uncertainty estimation. Such example results are presented in Figure 4b. The uncertainty maps exhibit high uncertainty regions around object boundaries and in regions containing complex textures. The far-right column depicts the pixels in pink for which the actual pixel intensity is not included in the corresponding credible interval. The fraction of pixels inside the 2.5%–97.5% credible interval is 95.34% (*Cameraman*), 98.28% (*House*), and 96.57% (*Barbara*). While the 95% chosen to define the credible intervals in Figure 4b does not directly relate to these numbers, the result in Figure 5a shows that the fraction of pixels inside the credible intervals drops nearly linearly as the width of the credible intervals shrinks, which indicates that the proposed EP method does not, on average across all the pixels, drastically underestimate or overestimate the marginal uncertainties in these cases.

The hyperparameters  $\{m_0, s^2\}$  have been estimated by EP-EM while running the experiments presented in the table of Figure 4a. Since the image intensity was already scaled in  $[0, 1]$ , the scaling factor  $\alpha$  has been fixed to 1 here. For completeness, EP-EM has been tested using the three images with 200 noise realizations for each standard deviation  $\sigma$ , resulting in 200 estimates ( $\hat{m}_0, \hat{s}^2$ ) per expert. In practice, the results obtained by different experts are similar, and we only present the results of one expert, selected randomly. The average ratios  $\hat{m}_0/m_0^{true}$  and  $\hat{s}^2/s_{true}^2$  and the associated 5th–95th quantile intervals are displayed in Figure 6. It can be seen that the 90% confidence intervals of  $\mathbb{E}[\hat{m}_0/m_0^{true}]$  and  $\mathbb{E}[\hat{s}^2/s_{true}^2]$  become larger as the noise variance increases, indicating that less information becomes available. However, all the experts provide similar results as they all use the same data. Consequently, although in Algorithm 3 we suggest using EP-EM for each expert, if communication is possible between experts, a single expert can perform the hyperparameter estimation and pass on its final estimates to the others, which can then only use Algorithm 1 for image restoration.

For denoising tasks with images of size  $256 \times 256$  pixels, the execution time of the Afonso





**Figure 6.** Results of hyperparameter estimation in denoising for Gaussian *i.i.d.* noise.  $\hat{m}_0$ ,  $\hat{s}^2$  are the estimates obtained by the proposed EP-EM method.  $\mathbb{E}[\hat{m}/m^{true}]$  and  $\mathbb{E}[\hat{s}^2/s_{true}^2]$  of the Y-axis are the 90% confidence intervals of the expected value of the estimated values divided by the ground truth values of three test images.

**Table 1**

Results of joint inpainting and denoising with Gaussian *i.i.d.* noise. PSNR (dB) comparison of different methods. The fraction of missing pixels is {40%, 60%, 80%}, with noise standard deviation  $\sigma = \{10, 20, 30\}/255$ .

	$\sigma$	Cameraman (256×256)			House (256×256)			Barbara (256×256)		
		40%	60%	80%	40%	60%	80%	40%	60%	80%
Afonso et al.	10/255	29.11	26.94	23.90	33.65	32.04	29.06	30.74	28.02	<b>24.33</b>
MAP-EPLL		29.05	27.14	24.13	33.70	31.97	28.61	29.92	26.92	23.46
Proposed		<b>29.74</b>	<b>27.56</b>	<b>24.24</b>	<b>34.21</b>	<b>32.75</b>	<b>29.38</b>	<b>31.09</b>	<b>28.27</b>	24.18
Afonso et al.	20/255	27.48	25.75	23.08	31.07	29.68	27.03	27.33	<b>25.58</b>	<b>23.29</b>
MAP-EPLL		26.74	25.05	22.28	30.27	28.96	25.49	24.98	23.36	21.52
Proposed		<b>27.85</b>	<b>26.11</b>	<b>23.41</b>	<b>31.36</b>	<b>30.05</b>	<b>27.15</b>	<b>27.38</b>	25.52	23.12
Afonso et al.	30/255	26.04	24.45	22.30	29.22	27.56	25.34	<b>25.20</b>	<b>24.04</b>	<b>22.42</b>
MAP-EPLL		24.72	23.09	20.66	28.24	31.88	23.18	22.86	21.87	20.42
Proposed		<b>26.24</b>	<b>24.60</b>	<b>22.59</b>	<b>29.26</b>	<b>27.61</b>	<b>25.37</b>	25.18	23.93	22.31

et al. method when taking non-overlapping patches is about 0.3 seconds, while each expert in the proposed EP-EM algorithm without parallel implementation takes 10.8 seconds on average (the algorithm converges in 2 to 3 iterations, and each iteration takes about 5 seconds). The costly computation is the grid search in finding the optimized parameters of  $m_0$  and  $s^2$ . For only EP with fixed hyperparameters, it takes an execution time similar to that of the method of Afonso et al. LIDIA takes about 9 seconds, BM3D takes about 1.6 seconds, and MAP-EPLL takes about 19.1 seconds.

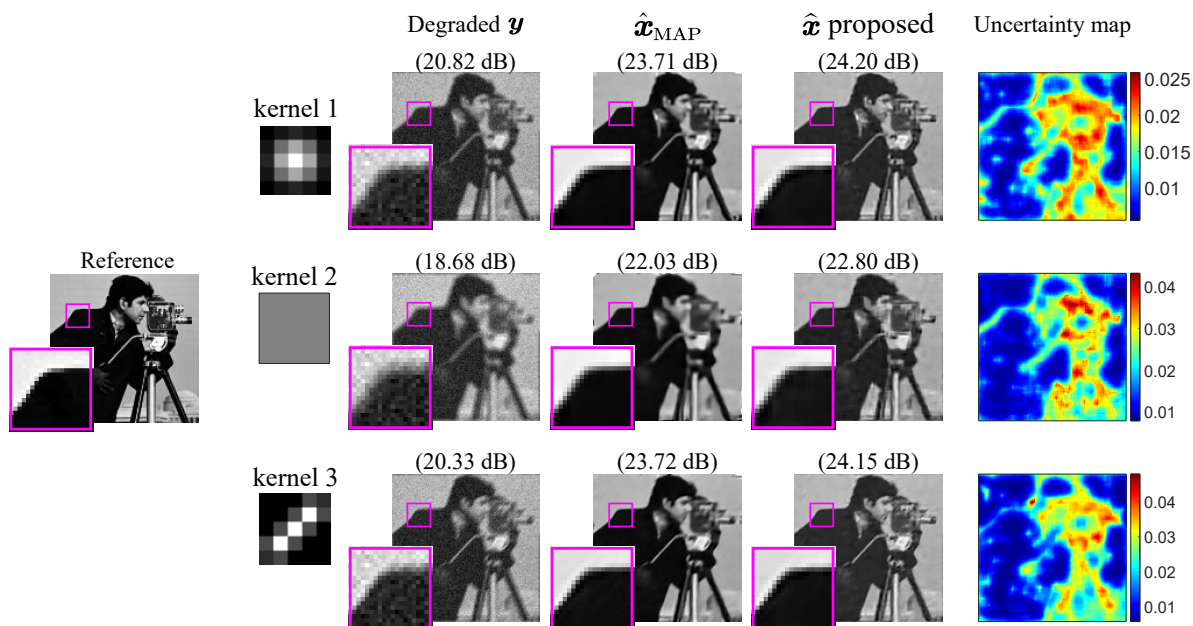
To assess the performance of the methods for inpainting problems, observed images  $\mathbf{y}$  are simulated via  $\mathbf{y} \sim \mathcal{N}(\mathbf{y}; \mathbf{H}\mathbf{x}, \sigma^2\mathbf{I})$ , where  $\mathbf{H}$  is a diagonal matrix with 40%, 60%, 80% diagonal elements randomly being zero, respectively, and  $\sigma = \{10, 20, 30\}/255$ . Table 1 reports the PSNRs of the inpainting results obtained using the proposed algorithm and two competing methods in [5] and [8]. It can be seen that here, the proposed EP-EM algorithm overall provides higher PSNR values than the MMSE patch-based GMM method [8] and MAP patch-based GMM method [5].

For these inpainting experiments, MAP-EPLL takes about 376.3 seconds, and the method

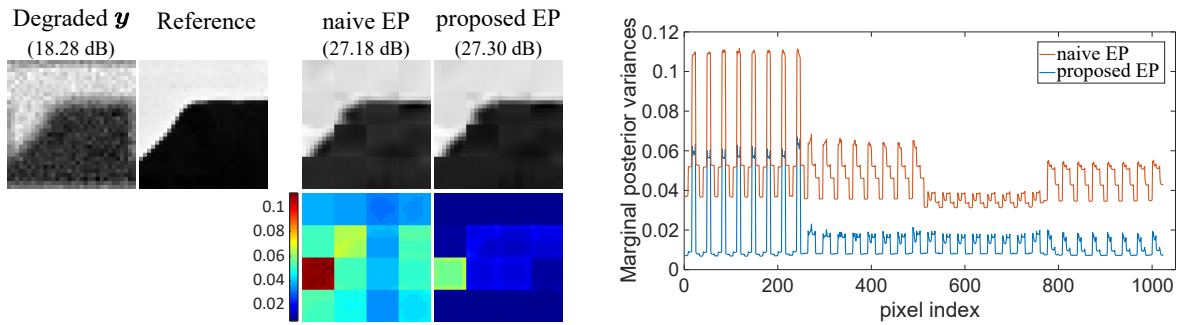
of Afonso et al. with only non-overlapping patches takes about 34.6 seconds, while a single expert of the proposed EP-EM method takes about 204.5 seconds on average (the algorithm converges in 7 iterations, and each iteration takes about 33.1 seconds). The hyperparameter estimation takes about 4.5 seconds at each iteration.

**6.1.2. Deconvolution.** For deconvolution experiments, a subimage of  $128 \times 128$  pixels taken from the *Cameraman* image is used as ground truth  $\mathbf{x}$  and convolved by three different kernels of size  $5 \times 5$  pixels, as shown in Figure 7 (i.e., a Gaussian blur, a uniform blur, and a motion blur). The proposed EP-EM method is compared to the MAP estimator MAP-EPLL [5]. It can be seen that the PSNR values after deconvolution using the proposed EP-EM method are higher than the MAP estimator. The zoomed-in boxes highlight the boundary between distinct homogeneous regions. The four uncertainty maps in the last column show that the pixels with high uncertainty generally appear at the edges of objects, while subtle changes can be observed when using different blur kernels.

For deconvolution of images of size  $128 \times 128$  pixels, the execution time of MAP-EPLL is about 7.0 seconds, while one expert of the proposed EP-EM methods takes about 750.1 seconds for 20 iterations (each iteration takes about 38.0 seconds, and no parallel implementation was used. 20 samples are generated in RBMC to approximate  $\text{Cov}_{P_1}(\mathbf{x})$  in (2.16)).



**Figure 7.** Results for deconvolution with Gaussian i.i.d. noise,  $\sigma = 0.05$ . First column: Reference of ground truth image and a zoomed-in region. Second column: Three kernels used in simulation. The size of all three kernels is  $5 \times 5$  pixels. Kernel 2 is a uniform kernel.  $\mathbf{H}$  is non-invertible constructed by kernel 1, and invertible constructed by kernels 2 and 3. Third column: Degraded observations with different kernels. Fourth column: Estimated images by MAP-EPLL [5]. Fifth column: Estimated images by the proposed EP-EM method. Last column: Uncertainty maps for the estimated images obtained by the proposed method.



(a) Approximate MMSE and associated uncertainty maps obtained by naive EP and the proposed EP algorithms.

(b) Marginal posterior variances obtained by naive EP and the proposed EP algorithms.

**Figure 8.** Comparison of approximate MMSE and posterior uncertainty quantification by naive EP and the proposed EP algorithms.

**6.1.3. Comparison with EP with full covariance matrices.** This subsection assesses the performance degradation of both the approximate MMSE estimates and the posterior uncertainty measures obtained by the proposed EP algorithm (Algorithm 1), when compared with more traditional EP implementation (referred to as naive EP), where all the approximating factors have full covariance matrices.

For image denoising and mild inpainting problems (with Gaussian noise), as discussed earlier, the exact posterior distribution for each expert is tractable and reduces to patch-wise independent GMMs. Thus our proposed EP approximation and the naive implementation both provide marginal means and variances that match those of the exact posterior distribution. For the deconvolution problem, the marginal moments estimated by EP are no longer exact. Within the naive implementation, all the covariance matrices of the approximating factors are full and the GMM patch-based prior is approximated by the product of  $J$  Gaussian approximating factors, where  $J$  is the number of non-overlapping patches in the image. This leads to  $(J + 1)$  factors that need to be updated sequentially in each EP iteration. To keep the computational cost of the naive implementation relatively low, we only consider a  $32 \times 32$  pixel image (a small subset of *Cameraman*), with  $\mathbf{H}$  corresponding to a  $5 \times 5$  pixel uniform kernel. Moreover, we only use a single expert in this comparison to better visualize the impact on a given posterior distribution. The results in Figure 8 show that the approximate MMSE estimates obtained by the naive EP and the proposed EP methods are comparable. In this example, it seems that the marginal variances of the naive EP method are larger than those of the proposed EP method. While we think that this is due to the fact that the block-diagonal approximation neglects part of the actual correlations, it could also be that the naive EP implementation overestimates uncertainties due to convergence issues during the gradient-based iterative updates used to ensure that all the approximate covariance matrices are strictly positive-definite. Indeed, the gradient descent used to minimize (2.13) might stop when moving in the gradient direction is no longer possible due to the constraints, but it might be possible to further optimize the cost function by moving along another direction (not investigated in this work). Note that the blocky artifacts in Figure 8 are due to the fact

PV	VST-BM3D	SPIRAL-TV	PIDAL-TV	Proposed	Degraded $\mathbf{y}$	$\hat{\mathbf{x}}$ proposed	Uncertainty map	<span style="color: magenta;">■</span> 1 <span style="color: white;">□</span> 0
Cameraman (256 × 256)								
10	<b>26.34</b>	24.01	24.00	<u>25.08</u>				
30	<b>28.57</b>	26.30	26.57	<u>27.77</u>				
50	<b>29.75</b>	27.76	27.93	<u>29.13</u>				
100	<b>31.27</b>	29.61	29.66	<u>30.95</u>				
255	<b>33.76</b>	32.08	32.38	<u>33.63</u>				
House (256 × 256)								
10	<b>28.47</b>	25.62	25.71	<u>26.25</u>				
30	<b>31.24</b>	28.47	28.41	<u>29.54</u>				
50	<b>32.50</b>	29.61	29.61	<u>31.11</u>				
100	<b>33.75</b>	31.21	31.18	<u>32.75</u>				
255	<b>35.80</b>	33.39	33.34	<u>34.85</u>				
Barbara (256 × 256)								
10	<b>25.17</b>	21.76	21.76	<u>23.18</u>				
30	<b>28.08</b>	24.08	24.09	<u>26.11</u>				
50	<b>29.41</b>	25.21	25.24	<u>27.57</u>				
100	<b>31.31</b>	27.10	27.13	<u>29.83</u>				
255	<b>33.62</b>	29.86	29.91	<u>32.66</u>				

(a) PSNR (dB) comparison.


(b) Images denoised by the proposed EP-EM algorithm (peak value = 30).

**Figure 9.** Denoising results for Poisson noise. (a) PSNR (dB) comparison of different denoising methods. The highest and second highest PSNR values are in bold and underlined, respectively. (b) Restored images obtained by the proposed EP-EM (second column), associated uncertainty maps (third column), and binary maps indicating pixels for which ground truth intensity falls into a 2.5%–97.5% credible interval (last column). Pixels are depicted in white (“0”) if the intensity falls into the credible region and in pink (“1”) otherwise.

that a only single expert is considered here. In terms of computational cost, the proposed EP algorithm only takes about 1.6 minutes for 20 iterations, while the naive EP method takes about 2 hours to obtain the posterior approximations (for the same 20 EP iterations). In addition, the naive EP cannot scale to large images as it requires the inversion of large  $N \times N$  covariance matrices, while the proposed EP algorithm allows for greater scalability, inverting only small  $r \times r$  covariance matrices.

**6.2. Poisson noise model.** This subsection demonstrates the effectiveness of the proposed patch-based EP algorithm when the observation noise is being Poisson distributed. The proposed method is compared with VST-BM3D [59], SPIRAL-TV [60], and PIDAL-TV [61]. VST-BM3D combines the Anscombe transform [62] with the BM3D algorithm. The regularization parameters of SPIRAL-TV and PIDAL-TV are tuned to find the highest PSNRs.

**6.2.1. Denoising and mild inpainting.** The same benchmark grayscale images (*Cameraman*, *House*, *Barbara*) as in the previous experiments are used in this section. For the denoising experiments, the observed images  $\mathbf{y}$  are simulated using  $\mathbf{y} \sim \mathcal{P}(\mathbf{x})$ , and the peak value of  $\mathbf{x}$  is set to be in the set of  $\{10, 30, 50, 100, 255\}$ . The table in Figure 9a presents the PSNR values obtained by different denoising methods. VST-BM3D provides higher PSNRs than the proposed EP algorithm, in particular when the peak value is low. This is mainly because in a low peak value regime, approximating the Poisson distribution of observations using a Gaussian distribution is difficult, and the accuracy loss may be substantial. Compared to SPIRAL-TV and PIDAL-TV, which are non-patch-based methods, the proposed EP



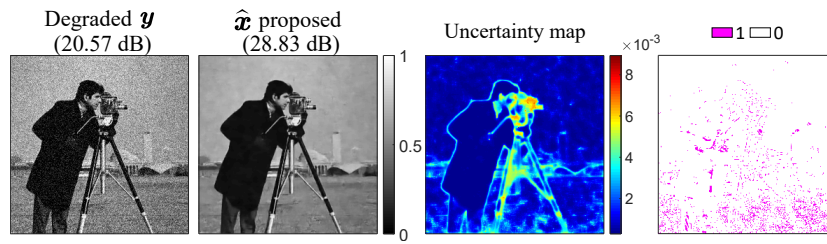

**Figure 10.** Results for hyperparameter estimation in denoising for Poisson noise.  $\hat{m}_0$ ,  $\hat{s}^2$ ,  $\hat{\alpha}$  are the estimates obtained by the proposed EP-EM method. Y-axis  $\mathbb{E}[\hat{m}/m^{true}]$ ,  $\mathbb{E}[\hat{s}^2/s_{true}^2]$ ,  $\mathbb{E}[\hat{\alpha}^2/\alpha^{true}]$  are the 90% confidence intervals of the expected value of the estimated values divided by the ground truth values of three test images.

algorithm performs better, and it does not require as much supervision in hyperparameter tuning as the two methods, since a subset of parameters in the GMM patch-based prior is pre-trained, and another subset of hyperparameters is retrieved in an unsupervised manner. Similar to denoising experiments for Gaussian i.i.d. noise, we emphasize that the main goal of the proposed EP algorithm is not to achieve the highest PSNRs. Rather, it aims to be scalable to high-dimensional imaging problems to provide comparable denoising results with additional uncertainty estimation. Figure 9b shows the uncertainty maps and binary images associated with uncertainty quantification for a peak value of 30. The fraction of pixels fall into 2.5%–97.5% credible interval is 92.89% (*Cameraman*), 98.43% (*House*), and 95.38% (*Barbara*). Figure 5b presents the fraction of pixels inside different credible intervals, which follows the same linear trend as in the denoising experiments with Gaussian noise.

The hyperparameter estimation results for the three images are presented in Figure 10. This figure shows that 90% confidence intervals of the three expected values  $\mathbb{E}[\hat{m}_0/m_0^{true}]$ ,  $\mathbb{E}[\hat{s}^2/s_{true}^2]$ ,  $\mathbb{E}[\hat{\alpha}^2/\alpha_{true}]$  are getting narrower as the peak value increases, indicating that as the peak value of the Poisson observations increases, it becomes easier to estimate the hyperparameters, as the signal-to-noise ratio increases.

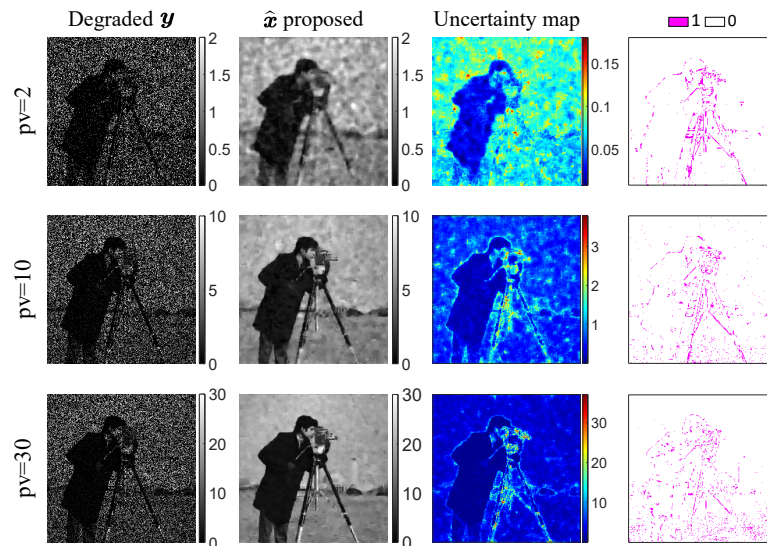
Figure 11 depicts a comparison of denoised *Cameraman* images corrupted by Gaussian i.i.d. noise (Figure 11a) and Poisson noise (Figure 11b) using the proposed EP-EM algorithm. It can be seen that the uncertainty maps present similar yet different features. Both maps exhibit lower uncertainties in homogeneous regions; however, the uncertainties in the bottom image (Poisson noise) are lower in homogeneous low-intensity regions (e.g., lower in the *Cameraman*'s coat and hat than in the sky). This is mainly due to two reasons: the Poisson likelihood which induces uncertainties depending on the mean signal values, and the Poisson rectified model which tends to keep most of the mass of  $Q(\mathbf{x})$  in the positive orthant.

We also conducted inpainting experiments, where the observed images  $\mathbf{y}$  are simulated using  $\mathbf{y} \sim \mathcal{P}(\mathbf{H}\mathbf{x})$ , the degradation operator  $\mathbf{H}$  is a diagonal matrix with 60% diagonal elements randomly being zero, and the peak value of  $\mathbf{x}$  is set to be 2, 10, and 30, respectively. Restoration results are shown in Figure 12. With the same percentage of missing pixels, the quality of the restored image and associated uncertainty map increases as the peak value increases. It is however worth noting that even when the peak value is very low (peak value = 2) and a large fraction of pixels is missing (60% missing), the proposed algorithm manages to restore a

(a) Denoising for Gaussian i.i.d. noise ( $\sigma = 25/255$ ).

(b) Denoising for Poisson noise (peak value = 30).

**Figure 11.** Comparison results between Gaussian and Poisson denoising. The fraction of pixels in the 2.5%–97.5% credible interval is (a) 95.34% and (b) 92.84%.



**Figure 12.** Comparison results of inpainting and denoising with Poisson noise. The percentage of missing pixels is 60%. The PSNR of the estimated images and the fraction of pixels that falls into the 2.5%–97.5% credible interval are 19.11 dB, 96.71% (peak value = 2), 22.14 dB, 96.51% (peak value = 10), and 24.26 dB, 96.25% (peak value = 30), respectively.

recognizable image, and the uncertainty map also presents low uncertainty in the flat region (the body of *Cameraman*). In addition, most pixels that are not in the 95% credible intervals construct the edges of objects.

For these denoising tasks, VST-BM3D takes about 0.3 seconds, SPIRAL-TV takes about



**Figure 13.** Results for deconvolution with Poisson noise. First column: Reference of ground truth image and a zoomed-in region. Second column: Observations degraded by a motion kernel of size  $5 \times 5$  pixels in 45 degree direction and Poisson noise with peak values  $\{10, 30, 50, 100, 255\}$ . Third column: Estimated images by VST-BM3D [59]. Fourth column: Estimated images by SPIRAL-TV [60]. Fifth column: Estimated images by PIDAL [61]. Sixth column: Estimated images by the proposed EP-EM method. Last column: Uncertainty maps for the estimated images obtained by the proposed method.

1.7 seconds, PIDAL-TV takes about 1.5 seconds, and one expert of the proposed EP method takes about 513.3 seconds for 13 iterations (each iteration takes about 44 seconds). In Algorithm 2, the main cost of computation in each iteration includes the following: (i) line 4, the 1D integrals in (3.7) for  $n = 1, \dots, 256 \times 256$  pixels in serial, (ii) lines 8-11, EP update of  $q_{x,1}(\mathbf{x})$ , and (iii) lines 16-19, EP update of  $q_{x,0}(\mathbf{x})$ .

**6.2.2. Deconvolution.** The same *Camerman* grayscale image as that used for deconvolution experiments with Gaussian noise is used here, but it is multiplied by peak values  $\{10, 30, 50, 100, 255\}$  to obtain the ground truth  $\mathbf{x}$ . The observed image  $\mathbf{y}$  is simulated by  $\mathbf{y} \sim \mathcal{P}(\mathbf{H}\mathbf{x})$ , and  $\mathbf{H}$  is formed by a  $5 \times 5$  pixel motion kernel with an angle of 45 degrees. Figure 13 presents examples of images restored by different methods. Although the proposed EP-EM method does not achieve the highest PSNRs, it does not produce artifacts as severe as VST-BM3D, SPIRAL-TV, or PIDAL-TV. In addition, it can be seen that the overall uncertainty of the restored images increases as the peak value increases, which is due to the fact

that as the mean of a Poisson distribution increases it induces larger posterior uncertainties, which are not completely reduced by the prior models.

For deconvolution of images of size  $128 \times 128$  pixels, VST-BM3D takes about 0.4 seconds, SPIRAL-TV takes about 3.7 seconds, PIDAL-TV takes about 13.8 seconds, and one expert of the proposed method takes about 869.3 seconds for 20 iterations (each iteration takes about 44 seconds on average). 20 samples are generated in RBMC to compute  $\text{Cov}_{P_1}(\mathbf{x})$  in (3.10)). The costly computation is similar to deconvolution with Gaussian i.i.d. noise; i.e., most of the computation can be reduced by parallel implementation.

In contrast to the Gaussian noise case, implementing a more traditional EP algorithm is not straightforward with Poisson noise. Our method using auxiliary variables benefits from a cavity distribution with a diagonal covariance matrix when the tilted distribution is associated with the Poisson likelihood. This allows the parallel computation of the resulting marginal moments. Without auxiliary variables and assuming full covariance matrices for the approximating factors, the resulting EP updates would need to be performed sequentially, using rank-1 updates. In our experience, such updates can become numerically unstable, especially when enforcing the covariance matrices to be positive definite. Thus, we did not include comparisons with alternative EP implementations with less restrictive covariance structure constraints.

**7. Conclusions and discussions.** In this paper, we proposed a new EP framework for image restoration using patch-based priors. In a similar fashion to VB, EP approximates the posterior distribution of interest by a simpler distribution whose moments are tractable. To address the computational challenges associated with large covariance matrix handling, we proposed considering a combination of isotropic, diagonal, and block-diagonal matrices, allowing, at least partially, distributed computation. Using augmented Bayesian models and factor graphs, we showed that the proposed method can be easily adapted to complex noise models. Indeed, for independently distributed noise realizations, the corresponding EP update simply requires the estimation of the first and second moments of 1D distributions, which can be performed in a parallel fashion. Moreover, if the initial GMM trained on external images whose scale does not match the current image of interest, it is possible to adjust some of its parameters (e.g., scale and offset) using a variational EM approach, which does not require crucial user supervision (in contrast to most MAP approaches). Although the proposed EP algorithm does not outperform state-of-the-art methods tailored for each restoration task in terms of quality of the estimated images, as an approximate MMSE estimator, it provides competitive posterior mean estimates with additional pixel-wise posterior uncertainty quantification, which holds great promise for uncertainty quantification in large scale imaging problems. Although we did not implement fully optimized versions of the EP-EM methods, the reported computational time can be further reduced by using a parallel update of blocks in covariance matrices of each expert, and parallel running for all the experts. This would however require significant additional work to optimize the updates for the different scenarios considered, and this is left to future work.

In practice, we observed that the Product-of-Experts approach provides very good results, and in particular, it allows blocky artifacts to be removed. Moreover, the resulting marginal uncertainties are in practice good indicators of how far the posterior mean is from the actual



pixel value. Hence, this information could be used for subsequent decision processes. It should however be noted that due to their approximate nature, such measures of uncertainty should be exploited with care. Since the exact posterior moments are however often not accessible, assessing the accuracy of the EP approximations remains an open problem.

Here, the block-diagonal covariance structure is motivated by the fact that the exact posterior covariance matrix is diagonally dominant. For deconvolution with Gaussian noise, we illustrated that the block-diagonal assumption does not significantly affect the estimated posterior means but it can modify the estimated marginal variances. Nonetheless, such structural constraints seemed necessary for EP to be scalable in the Poisson regime. If the blur kernel is larger than the patches considered, we observed that the reconstruction performance and convergence speed of EP can be degraded. Similarly, if the operator  $\mathbf{H}$  introduces strong far-field correlation between pixels, as can happen in compressive sensing or tomography problems and when global image prior models are used, the method proposed in this paper does not provide satisfactory results. In such cases, different structural constraints should be investigated to make sure the EP method remains scalable and flexible enough to capture partially posterior correlations. While the EP covariance structures can be informed by the correlations expected to be present in the exact posterior distribution, they should also be designed and exploited together with the structures of the EP factor graphs, in particular when using data augmentation. For instance, even if  $Q(\mathbf{x})$  is designed to have a block-diagonal covariance matrix, not all the approximating factors need to satisfy the same constraints.

**Appendix A. Details of KL divergence minimization with different covariance structures.** When minimizing the KL divergence in the form of  $KL(P||Q)$ , where  $P$  is a tilted distribution and  $Q$  is an EP Gaussian approximation, the divergence can be parameterized by the mean and covariance matrix of  $Q$ . Here, for simplicity,  $P$  and  $Q$  can represent densities of  $\mathbf{u}$  or  $\mathbf{x}$ . In the following derivation, we use the following general notation for the mean and covariance matrices, as listed in Table 2. Moreover, we have  $\boldsymbol{\Omega}_* = \boldsymbol{\Omega}_0 + \boldsymbol{\Omega}_1$ ,  $\boldsymbol{\Omega}_*\boldsymbol{\mu}_* = \boldsymbol{\Omega}_0\boldsymbol{\mu}_0 + \boldsymbol{\Omega}_1\boldsymbol{\mu}_1$ .

Table 2

Notation for the mean and covariance matrices of  $P$ ,  $Q$ , and individual approximating factors  $q_0(\cdot)$ ,  $q_1(\cdot)$ .

	$P(\cdot)$	$Q(\cdot)$	$q_0(\cdot)$	$q_1(\cdot)$
Mean	$\mathbb{E}_P[\cdot]$	$\boldsymbol{\mu}_*$	$\boldsymbol{\mu}_0$	$\boldsymbol{\mu}_1$
Covariance matrix	$\text{Cov}_P(\cdot)$	$\boldsymbol{\Sigma}_*$	$\boldsymbol{\Sigma}_0$	$\boldsymbol{\Sigma}_1$
Precision matrix		$\boldsymbol{\Omega}_*$	$\boldsymbol{\Omega}_0$	$\boldsymbol{\Omega}_1$

As explained in section 2, if the mean of  $Q$  is unconstrained, the KL divergence is minimized when  $\boldsymbol{\mu}_* = \mathbb{E}_P[\cdot]$ , and we only need to optimize  $\boldsymbol{\Omega}_*$ . The loss function reduces to

$$(A.1) \quad F(\boldsymbol{\Omega}_i) \propto -\log(\det(\boldsymbol{\Omega}_i + \boldsymbol{\Omega}_{\setminus i})) + \langle \boldsymbol{\Omega}_i + \boldsymbol{\Omega}_{\setminus i}, \text{Cov}_P(\cdot) \rangle,$$

which is convex w.r.t.  $\boldsymbol{\Omega}_i \in \mathbb{S}_{++}^N$ . Therefore, without additional constraints,  $\boldsymbol{\Omega}_i$  can be obtained by solving

$$(A.2) \quad \hat{\boldsymbol{\Omega}}_i = \underset{\boldsymbol{\Omega}_i \in \mathbb{S}_{++}^N}{\text{argmin}} -\log(\det(\boldsymbol{\Omega}_i + \boldsymbol{\Omega}_{\setminus i})) + \langle \boldsymbol{\Omega}_i + \boldsymbol{\Omega}_{\setminus i}, \text{Cov}_P(\cdot) \rangle.$$

**A.1. EP update of diagonal covariance.** In section 2, we discussed the update of  $\mathbf{\Omega}_i$ , assuming that both  $\mathbf{\Omega}_i$  and  $\mathbf{\Omega}_{\setminus i}$  are block-diagonal. Here, we consider the case where both  $\mathbf{\Omega}_i$  and  $\mathbf{\Omega}_{\setminus i}$  are diagonal while optimizing  $\mathbf{\Omega}_i$ . For brevity, we use  $[d_{i,1}, \dots, d_{i,N}] = \text{diag}(\text{Cov}_{P_i}(\cdot))$ ,  $[p_{i,1}, \dots, p_{i,N}] = \text{diag}(\mathbf{\Omega}_i)$ , and  $[p_{\setminus i,1}, \dots, p_{\setminus i,N}] = \text{diag}(\mathbf{\Omega}_{\setminus i})$ . In that case, (A.1) becomes

$$(A.3) \quad F(\mathbf{\Omega}_i) \propto \sum_{n=1}^N -\log(p_{i,n} + p_{\setminus i,n}) + (p_{i,n} + p_{\setminus i,n})d_{i,n},$$

which can be optimized independently w.r.t. the elements of  $\text{diag}(\mathbf{\Omega}_i)$ . For the element  $p_{i,n}$ , the cost function becomes

$$(A.4) \quad F(p_{i,n}) \propto -\log(p_{i,n} + p_{\setminus i,n}) + (p_{i,n} + p_{\setminus i,n})d_{i,n}.$$

This function is convex w.r.t.  $p_{i,n}$  and has a unique minimizer,  $p_{i,n} = \frac{1}{d_{i,n}} - p_{\setminus i,n}$ . If this minimizer is negative, it is replaced by a small positive value (e.g.,  $10^{-8}$ ), leading to a large variance in  $\Sigma_i$  for that element.

**A.2. EP update of isotropic covariance.** Based on (A.3), when the diagonal elements of the precision matrix  $\mathbf{\Omega}_i$  are constrained to be equal to a same value denoted  $p$ , the loss function becomes

$$(A.5) \quad F(\mathbf{\Omega}_i) \propto \sum_{n=1}^N -\log(p + p_{\setminus i,n}) + (p + p_{\setminus i,n})d_{i,n}.$$

This function is convex w.r.t.  $p$ , and it can be minimized using a Newton–Raphson method. At the  $(t)$ th iteration, the update is

$$(A.6) \quad p_{(t)} = \max \left( p_{(t-1)} + \frac{\sum_{n=1}^N \frac{1}{p_{(t-1)} + p_{\setminus i,n}} - \sum_{n=1}^N d_{i,n}}{\sum_{n=1}^N \frac{1}{(p_{(t-1)} + p_{\setminus i,n})^2}}, \epsilon \right),$$

where  $\epsilon = 10^{-8}$  is used to ensure that the covariance matrix remains positive definite.

## REFERENCES

- [1] M. A. T. FIGUEIREDO AND R. D. NOWAK, *An EM algorithm for wavelet-based image restoration*, IEEE Trans. Image Process., 12 (2003), pp. 906–916.
- [2] F. ABRAMOVICH, T. SAPATINAS, AND B. W. SILVERMAN, *Wavelet thresholding via a Bayesian approach*, J. R. Stat. Soc. Ser. B Stat. Methodol., 60 (1998), pp. 725–749.
- [3] J. M. BIOUSCAS-DIAS, *Bayesian wavelet-based image deconvolution: A GEM algorithm exploiting a class of heavy-tailed priors*, IEEE Trans. Image Process., 15 (2006), pp. 937–951.
- [4] H. RUE AND L. HELD, *Gaussian Markov Random Fields: Theory and Applications*, CRC Press, Boca Raton, FL, 2005.
- [5] D. ZORAN AND Y. WEISS, *From learning models of natural image patches to whole image restoration*, Proceedings of the IEEE International Conference on Computer Vision (ICCV), IEEE, Washington, DC, 2011, pp. 479–486.

- [6] M. NIKNEJAD, J. BIOUCAS-DIAS, AND M. A. T. FIGUEIREDO, *External patch-based image restoration using importance sampling*, IEEE Trans. Image Process., 28 (2019), pp. 4460–4470.
- [7] M. ELAD, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*, Springer, New York, 2010.
- [8] A. M. TEODORO, M. S. C. ALMEIDA, AND M. A. T. FIGUEIREDO, *Single-frame image denoising and inpainting using Gaussian mixtures*, in Proceedings of ICPRAM (2), Lisbon, Portugal, 2015, pp. 283–288.
- [9] M. NIKNEJAD, H. RABBANI, AND M. BABAIE-ZADEH, *Image restoration using Gaussian mixture models with spatially constrained patch clustering*, IEEE Trans. Image Process., 24 (2015), pp. 3624–3636.
- [10] A. M. TEODORO, J. M. BIOUCAS-DIAS, AND M. A. T. FIGUEIREDO, *Block-Gaussian-mixture priors for hyperspectral denoising and inpainting*, IEEE Trans. Geosci. Remote Sensing, 59 (2021), pp. 2478–2486.
- [11] R. ZHANG, D. H. YE, D. PAL, J.-B. THIBAUT, K. D. SAUER, AND C. A. BOUMAN, *A Gaussian mixture MRF for model-based iterative reconstruction with applications to low-dose X-ray CT*, IEEE Trans. Comput. Imaging, 2 (2016), pp. 359–374.
- [12] S. ROTH AND M. J. BLACK, *Fields of experts: A framework for learning image priors*, in Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), Vol. 2, IEEE, Washington, DC, 2005, pp. 860–867.
- [13] S. V. VENKATAKRISHNAN, C. A. BOUMAN, AND B. WOHLBERG, *Plug-and-play priors for model based reconstruction*, in Proceedings of the 2013 IEEE Global Conference on Signal and Information Processing, IEEE, Washington, DC, 2013, pp. 945–948.
- [14] Y. ROMANO, M. ELAD, AND P. MILANFAR, *The little engine that could: Regularization by denoising (RED)*, SIAM J. Imaging Sci., 10 (2017), pp. 1804–1844, <https://doi.org/10.1137/16M1102884>.
- [15] D. ULYANOV, A. VEDALDI, AND V. LEMPITSKY, *Deep image prior*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, IEEE, Washington, DC, 2018, pp. 9446–9454.
- [16] G. MATAEV, P. MILANFAR, AND M. ELAD, *DeepRED: Deep image prior powered by RED*, in Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops, IEEE, Washington, DC, 2019, pp. 0–0.
- [17] G. VAKSMAN, M. ELAD, AND P. MILANFAR, *LIDIA: Lightweight Learned Image Denoising with Instance Adaptation*, in IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, IEEE, Washington, DC, 2020, pp. 2220–2229.
- [18] K. ZHANG, W. ZUO, Y. CHEN, D. MENG, AND L. ZHANG, *Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising*, IEEE Trans. Image Process., 26 (2017), pp. 3142–3155.
- [19] J. M. BERNARDO AND A. F. M. SMITH, *Bayesian Theory*, Wiley Ser. Probab. Statist. Probab. Statist., John Wiley & Sons, Chichester, UK, 1994.
- [20] M. PEREYRA, *Maximum-a-posteriori estimation with Bayesian confidence regions*, SIAM J. Imaging Sci., 10 (2017), pp. 285–302, <https://doi.org/10.1137/16M1071249>.
- [21] A. REPETTI, M. PEREYRA, AND Y. WIAUX, *Scalable Bayesian uncertainty quantification in imaging inverse problems via convex optimization*, SIAM J. Imaging Sci., 12 (2019), pp. 87–118, <https://doi.org/10.1137/18M1173629>.
- [22] A. DURMUS, É. MOULINES, AND M. PEREYRA, *Efficient Bayesian computation by proximal Markov chain Monte Carlo: When Langevin meets Moreau*, SIAM J. Imaging Sci., 11 (2018), pp. 473–506, <https://doi.org/10.1137/16M1108340>.
- [23] C. M. BISHOP, *Pattern Recognition and Machine Learning*, Springer, New York, 2006.
- [24] D. M. BLEI, A. KUCUKELBIR, AND J. D. MCAULIFFE, *Variational inference: A review for statisticians*, J. Amer. Statist. Assoc., 112 (2017), pp. 859–877.
- [25] T. P. MINKA, *Expectation propagation for approximate Bayesian inference*, in Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, UAI’01, Seattle, WA, 2011, pp. 362–369.
- [26] M. OPPER AND O. WINTHER, *Gaussian processes for classification: Mean-field algorithms*, Neural Comput., 12 (2000), pp. 2655–2684.
- [27] A. GELMAN, A. VEHTARI, P. JYLÄNKI, C. ROBERT, N. CHOPIN, AND J. P. CUNNINGHAM, *Expectation propagation as a way of life*, J. Mach. Learn. Res., 21 (2020), pp. 1–53.
- [28] M. SEEGER AND H. NICKISCH, *Fast convergent algorithms for expectation propagation approximate Bayesian inference*, in Proceedings of the Fourteenth International Conference on Artificial Intel-

- ligence and Statistics, Ft. Lauderdale, FL, 2011, pp. 652–660.
- [29] M. SEEGER, F. STEINKE, AND K. TSUDA, *Bayesian inference and optimal design in the sparse linear model*, in Artificial Intelligence and Statistics, PMLR, 2007, pp. 444–451.
- [30] M. W. SEEGER AND H. NICKISCH, *Large scale Bayesian inference and experimental design for sparse linear models*, SIAM J. Imaging Sci., 4 (2011), pp. 166–199, <https://doi.org/10.1137/090758775>.
- [31] J. M. HERNÁNDEZ-LOBATO, D. HERNÁNDEZ-LOBATO, AND A. SUÁREZ, *Expectation propagation in linear regression models with spike-and-slab priors*, Mach. Learn., 99 (2015), pp. 437–487.
- [32] Y. ALTMANN, A. PERELLI, AND M. E. DAVIES, *Expectation-propagation algorithms for linear regression with Poisson noise: Application to photon-limited spectral unmixing*, in Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, Washington, DC, 2019, pp. 5067–5071.
- [33] A. S. I. KIM AND M. P. WAND, *On expectation propagation for generalised, linear and mixed models*, Aust. N. Z. J. Stat., 60 (2018), pp. 75–102.
- [34] N. HEES, D. TARLOW, AND J. WINN, *Learning to pass expectation propagation messages*, in Advances in Neural Information Processing Systems, Vol. 26, Curran Associates, Red Hook, NY, 2013, pp. 3219–3227.
- [35] A. BRAUNSTEIN, A. MUNTONI, A. PAGNANI, AND M. PIEROPAN, *Compressed sensing reconstruction using expectation propagation*, J. Phys. A, 53 (2019), 184001.
- [36] A. P. MUNTONI, R. D. H. ROJAS, A. BRAUNSTEIN, A. PAGNANI, AND I. P. CASTILLO, *Nonconvex image reconstruction via expectation propagation*, Phys. Rev. E, 100 (2019), 032134.
- [37] A. BRAUNSTEIN, A. P. MUNTONI, AND A. PAGNANI, *An analytic approximation of the feasible space of metabolic networks*, Nature Commun., 8 (2017), 14915.
- [38] C. ZHANG, S. ARRIDGE, AND B. JIN, *Expectation propagation for Poisson data*, Inverse Problems, 35 (2019), 085006.
- [39] J. P. CUNNINGHAM, P. HENNIG, AND S. LACOSTE-JULIEN, *Gaussian Probabilities and Expectation Propagation*, preprint, <https://arxiv.org/abs/1111.6832>, 2011.
- [40] B. POTETZ AND M. HAJIARBABI, *Whitened expectation propagation: Non-Lambertian shape from shading and shadow*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Washington, DC, 2013, pp. 1674–1681.
- [41] D. L. DONOHO, A. MALEKI, AND A. MONTANARI, *Message-passing algorithms for compressed sensing*, Proc. Natl. Acad. Sci. USA, 106 (2009), pp. 18914–18919.
- [42] S. RANGAN, *Generalized approximate message passing for estimation with random linear mixing*, in Proceedings of the IEEE International Symposium on Information Theory, IEEE, Washington, DC, 2011, pp. 2168–2172.
- [43] P. SCHNITER, S. RANGAN, AND A. K. FLETCHER, *Vector approximate message passing for the generalized linear model*, in Proceedings of the IEEE 50th Asilomar Conference on Signals, Systems and Computers, IEEE, Washington, DC, 2016, pp. 1525–1529.
- [44] J. BARZILAI AND J. M. BORWEIN, *Two-point step size gradient methods*, IMA J. Numer. Anal., 8 (1988), pp. 141–148.
- [45] M. R. HESTENES AND E. STIEFEL, *Methods of conjugate gradients for solving linear systems*, J. Research Nat. Bur. Standards, 49 (1952), pp. 409–436.
- [46] P. SIDÉN, F. LINDGREN, D. BOLIN, AND M. VILLANI, *Efficient covariance approximations for large sparse precision matrices*, J. Comput. Graph. Statist., 27 (2018), pp. 898–909.
- [47] Y.-J. KO AND M. W. SEEGER, *Expectation propagation for rectified linear Poisson regression*, in Proceedings of the Asian Conference on Machine Learning, Hamilton, Auckland, New Zealand, 2016, pp. 253–268.
- [48] Y. J. KO, *Applications of Approximate Learning and Inference for Probabilistic Models*, Tech. report, EPFL, Lausanne, Switzerland, 2017.
- [49] M. P. WAND, J. T. ORMEROD, S. A. PADOAN, AND R. FRÜHWIRTH, *Mean field variational Bayes for elaborate distributions*, Bayesian Anal., 6 (2011), pp. 847–900.
- [50] Y. ALTMANN, S. McLAUGHLIN, AND M. PADGETT, *Unsupervised restoration of subsampled images constructed from geometric and binomial data*, in Proceedings of the 7th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), IEEE, Washington, DC, 2017, pp. 1–5.

- [51] T. MINKA, *Divergence Measures and Message Passing*, Tech. report MSR-TR-2005-173, Microsoft Research, 2005.
- [52] A. GELMAN, J. B. CARLIN, H. S. STERN, D. B. DUNSON, A. VEHTARI, AND D. B. RUBIN, *Bayesian Data Analysis*, CRC Press, Boca Raton, FL, 2013.
- [53] G. E. HINTON, *Products of experts*, in Proceedings of the International Conference on Artificial Neural Networks (ICANN), IEEE, Washington, DC, 1999, pp. 1–6.
- [54] T. K. MOON, *The expectation-maximization algorithm*, IEEE Signal Process. Mag., 13 (1996), pp. 47–60.
- [55] G. CELEUX, F. FORBES, AND N. PEYRARD, *EM procedures using mean field-like approximations for Markov model-based image segmentation*, Pattern Recognition, 36 (2003), pp. 131–144.
- [56] H.-C. KIM AND Z. GHAHRAMANI, *Bayesian Gaussian process classification with the EM-EP algorithm*, IEEE Trans. Pattern Anal. Mach. Intell., 28 (2006), pp. 1948–1959.
- [57] D. MARTIN, C. FOWLKES, D. TAL, AND J. MALIK, *A Database of Human Segmented Natural Images and Its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics*, Research paper, Department of Electrical Engineering and Computer Sciences, University of California Berkeley, Berkeley, CA, 2001.
- [58] K. DABOV, A. FOI, V. KATKOVNIK, AND K. EGIAZARIAN, *Image denoising by sparse 3-D transform-domain collaborative filtering*, IEEE Trans. Image Process., 16 (2007), pp. 2080–2095.
- [59] M. MAKITALO AND A. FOI, *Optimal inversion of the Anscombe transformation in low-count Poisson image denoising*, IEEE Trans. Image Process., 20 (2010), pp. 99–109.
- [60] Z. T. HARMANY, R. F. MARCIA, AND R. M. WILLETT, *This is SPIRAL-TAP: Sparse Poisson intensity reconstruction algorithms—theory and practice*, IEEE Trans. Image Process., 21 (2011), pp. 1084–1096.
- [61] M. A. T. FIGUEIREDO AND J. M. BIOCAS-DIAS, *Restoration of Poissonian images using alternating direction optimization*, IEEE Trans. Image Process., 19 (2010), pp. 3133–3145.
- [62] F. J. ANSCOMBE, *The transformation of Poisson, binomial and negative-binomial data*, Biometrika, 35 (1948), pp. 246–254.