



A Hyperprior Bayesian Approach for Solving Image Inverse Problems

Cecilia Aguerrebere, Andrés Almansa, Yann Gousseau, Julie Delon, Pablo Musé

► To cite this version:

Cecilia Aguerrebere, Andrés Almansa, Yann Gousseau, Julie Delon, Pablo Musé. A Hyperprior Bayesian Approach for Solving Image Inverse Problems. 2014. hal-01107519v1

HAL Id: hal-01107519

<https://hal.science/hal-01107519v1>

Preprint submitted on 20 Jan 2015 (v1), last revised 15 May 2017 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Hyperprior Bayesian Approach for Solving Image Inverse Problems

Cecilia Aguerrebere
Duke University

cecilia.aguerrebere@duke.edu

Julie Delon
Université Paris Descartes

julie.delon@parisdescartes.fr

Andrés Almansa, Yann Gousseau
Télécom ParisTech

gousseau, almansa@telecom-paristech.fr

Pablo Musé
Universidad de la República, Uruguay

pmuse@fing.edu.uy

Abstract

Patch models have proven successful to solve a variety of inverse problems in image restoration. Recent methods, combining patch models with a Bayesian approach, achieve state-of-the-art results in several restoration problems. Different strategies are followed to determine the patch models, such as a fixed number of models to describe all image patches or a locally determined model for each patch. Local model estimation has proven very powerful for image denoising, but it becomes seriously ill-posed for other inverse problems such as interpolation of random missing pixels or zooming. In this work, we present a new framework for image restoration that combines these two powerful approaches: Bayesian restoration and a local characterization of image patches. By making use of a prior on the model parameters, we overcome the ill-posedness of the local estimation and obtain state-of-the-art results in problems such as interpolation, denoising and zooming. Experiments conducted on synthetic and real data show the effectiveness of the proposed approach.

1. Introduction

Digital images are subject to a wide variety of degradations, which in most cases can be modeled as transformations

$$\mathbf{Z} = \mathbf{D}\mathbf{C} + \mathbf{N}. \quad (1)$$

Different settings of the degradation matrix $\mathbf{D}$ model different problems such as zooming, deblurring or random missing pixels. Different characterizations of the noise term $\mathbf{N}$ describe noise degradations, ranging from the classical additive Gaussian noise to more complicated an realistic models such as multiplicative or signal dependent noise. These degradations are often combined in practice. For instance, raw images captured with regular digital cameras combine

signal dependent noise, limited spatial resolution and limited dynamic range, among others [2].

Inspired by the patch-based approach for texture synthesis proposed by Efros and Leung [7], Buades et al. [3] introduced the use of patches and the self-similarity hypothesis to the denoising problem leading to a new era of patch-based image restoration techniques. A major step forward in fully exploiting patches potential was taken with the introduction of patches prior models. Recent state-of-the-art methods make use of patch models in a Bayesian framework to restore degraded images. Some of them are devoted to the denoising problem [11, 5, 8, 15], while others propose a more general framework for the solution of image inverse problems [18, 16], including for instance inpainting, deblurring and zooming. The work by Lebrun et al. [10, 8] presents a thorough and very interesting analysis of several recent restoration methods, revealing their common roots and their relationship with the Bayesian approach.

Among the state-of-the-art restoration methods, two noticeable approaches are the patch-based Bayesian approach by Yu et al. [16], namely the piecewise linear estimators (PLE), and the non-local Bayes (NLB) algorithm by Lebrun et al. [8]. PLE is a general framework for the solution of image inverse problems under model (1), while NLB is a denoising method ($\mathbf{D} = \mathbf{Id}$). Both methods use a Gaussian patch prior learnt from image patches through iterative procedures. In the case of PLE, patches are modeled according to a Gaussian Mixture Model (GMM), with a relatively small number of classes (19 in all their experiments), whose parameters are learnt from all image patches¹. In the case of NLB, the parameters of the Gaussian model for each patch are computed from similar patches chosen from a local neighborhood, so the number of classes is not limited a priori (one class per patch). NLB outperforms PLE in the denoising task [14], mostly due to its continuous classifica-

¹Actually, the authors report the use of 128×128 image sub-regions in their experiments, so we may consider PLE as a semi-local approach.

tion model. Nevertheless, PLE obtains state-of-the-art results in other applications such as interpolation of missing pixels, deblurring and zooming. In particular, PLE yields very good results in various cases of interpolation of random missing pixels with high masking rates. A variant of PLE for inpainting is proposed by Wang [13] (E-PLE), using a GMM initialized from natural images instead of using synthetic images of edges as it is done in PLE.

Zoran and Weiss [18] (EPLL) follow a similar approach, but instead of iteratively updating the GMM from image patches, they use a larger number of classes (200) that are fixed and learnt from a large database of natural image patches (2×10^6 patches). Wang and Morel [15] claim that, in the case of denoising, it is better to have fewer models that are updated with the image patches (as in PLE) than having a large number of fixed models (as in EPLL). Moreover, unlike the previous methods, EPLL restores image patches according to the GMM prior² while keeping the restored image close to the corrupted image for a given corruption model.

In this work, we focus on the family of image degradation problems modeled by (1). We propose to take advantage of the proven restoration power of the Bayesian approach combined with a Gaussian patch prior [16], and boost it with the accuracy of local model estimation [8]. For this purpose, we propose to model image patches according to a Gaussian prior, whose parameters will be estimated locally from similar patches. The main challenge with this framework is to estimate the Gaussian parameters, i.e. the mean μ and the covariance matrix Σ , from a set of patches with potentially high degradation levels. For example, in the case of interpolation of random missing pixels with a masking rate of 70%, the patches used for the estimation of μ and Σ will lack 70% of the pixels, thus making the estimation problem very ill-posed. In order to tackle this problem, we include prior knowledge on the model parameters making use of what is known as an hyperprior, i.e. a probability distribution on the parameters of the prior. In Bayesian statistics, μ and Σ are known as hyperparameters, since they are the parameters of a prior distribution, while the prior on them is called an hyperprior. The use of an hyperprior allows to estimate μ and Σ from similar patches even if they present high degradation levels. The information provided by the hyperprior compensates for the patches missing information. Finally, image patches are restored using the maximum a posteriori (MAP) estimator with the computed Gaussian model.

The article is organized as follows. Section 2 presents the proposed approach and Section 3 presents the main implementation considerations. Section 4 presents supportive experiments and Section 5 summarizes the conclusions.

²EPLL does not impose a given prior, GMM is an option among others.

2. Hyperprior Bayesian Estimator

We describe here the proposed restoration method, Hyperprior Bayesian Estimator (HBE), that assumes a Gaussian prior for image patches whose parameters μ and Σ are estimated locally from a group of patches similar to the current patch. The proposed method iteratively alternates two steps. First, the patch model parameters μ and Σ are computed combining similar patches and the hyperprior on them. Then, the patch is restored computing its MAP estimate under the Gaussian prior defined by μ and Σ . Figure 1 shows a diagram of the proposed iterative approach.

2.1. Patch model

The observed image z is decomposed into I overlapping patches $z_{i=1,\dots,I}$ of size $\sqrt{n} \times \sqrt{n}$. Each patch $z_i \in \mathbb{R}^{n \times 1}$ is considered to be a realization of the random variable Z_i given by

$$Z_i = D_i C_i + N_i, \quad (2)$$

where $D_i \in \mathbb{R}^{n \times n}$ is a degradation operator, $C_i \in \mathbb{R}^{n \times 1}$ is the original patch we seek to estimate and $N_i \in \mathbb{R}^{n \times 1}$ is an additive noise term. The original patch C_i is modeled according to a Gaussian prior with mean μ_i and covariance matrix Σ_i . The noise term N_i depends on C_i , with its j -th entry given by $N_i^j = f(C_i^j) \varepsilon_i^j$, $j = 1, \dots, n$. Variables $(\varepsilon_i^j)_{j=1,\dots,n}$ are i.i.d Gaussian random variables with zero mean and unit variance, ε_i^j is independent of $(C_i^j)_{j=1,\dots,N}$ and $f : \mathbb{R} \rightarrow \mathbb{R}$ is the function describing the relationship between the noise variance and the signal value at each pixel. The entries of vector N_i are independent and its diagonal covariance matrix is denoted by Σ_{N_i} . The matrix D_i and the function f are assumed to be known. Model (2) includes the following degradation cases:

- D_i : resolution change (zooming), random missing pixels (inpainting)
- N_i : additive Gaussian noise with constant variance, variable variance and variance dependent on the pixel value (approximate Poisson noise), multiplicative noise.

The approximation of Poisson noise as additive Gaussian noise with variance depending on the pixel value is particularly useful in the case of raw images captured with regular digital cameras. In this case, the noise is independent among pixels but its variance depends on the irradiance reaching the pixel [2], the function f being

$$f(C_i^j) = (g^2 a_p \tau D_i^j C_i^j + \sigma_R^2) / (g a_p \tau)^2, \quad (3)$$

where g is the camera gain, a_p is the photo-response non uniformity factor, τ is the exposure time, μ_R and σ_R^2 are the readout noise mean and variance and $D_i^j = 1_{C_i^j < z_{sat}}$ models the saturation effect taking value one if the pixel is below the saturation threshold z_{sat} and zero otherwise [2].




Figure 1. Diagram of the proposed iterative approach with the corresponding initialization stage.

2.2. Patch restoration

Let us first suppose that the patch model parameters μ_i and Σ_i are known for each image patch z_i (its computation is explained in Section 2.3). We propose to restore each patch z_i by minimizing its Bayes risk under Model (2). In order to get a tractable, closed-form expression, we restrict the solution space to *affine estimators* only.

Proposition 1. *The affine estimator $\tilde{C}_i$ that minimizes the Bayes risk $\mathbb{E}[(\tilde{C}_i - C_i)^2]$ under Model (2) is given by*

$$\tilde{C}_i = \Sigma_i D_i^T (D_i \Sigma_i D_i^T + \Sigma_{N_i})^{-1} (Z_i - D_i \mu_i) + \mu_i. \quad (4)$$

The proof of Proposition 1 is presented in the supplementary material.

2.3. Estimation of μ_i and Σ_i

We propose to estimate μ_i and Σ_i by computing their MAP estimates using a set of patches similar to z_i , combined with a prior knowledge on μ_i and Σ_i , through an hyperprior. On the one hand, using similar patches gives a spatially adaptive or local characterization of the patch [8]. On the other hand, including the hyperprior makes the parameter estimation more robust, which is critical when few similar patches are available or some pixels are unknown (e.g. for interpolation and zooming). The dependence on the patch index i is hereafter omitted to simplify notation.

MAP estimates of μ and Σ To simplify calculations, it is convenient to work in terms of the precision matrix $\Lambda = \Sigma^{-1}$ instead of the covariance matrix Σ . As it is usual when considering hyperpriors, we rely on a conjugate distribution. In our case, that boils down to assuming a Normal-Wishart³ prior for the couple (μ, Λ) ,

$$\begin{aligned} p(\mu, \Lambda) &= \mathcal{N}(\mu | \mu_0, (\kappa \Lambda)^{-1}) \mathcal{W}(\Lambda | (\nu \Sigma_0)^{-1}, \nu) \\ &\propto |\Lambda|^{1/2} \exp\left(-\frac{\kappa}{2}(\mu - \mu_0)^T \Lambda (\mu - \mu_0)\right) \\ &\quad |\Lambda|^{(\nu-d-1)/2} \exp\left(-\frac{1}{2}\text{tr}(\nu \Sigma_0 \Lambda)\right), \end{aligned} \quad (5)$$

³The Normal-Wishart distribution is the conjugate prior of a multivariate normal distribution with unknown mean and covariance matrix. $\mathcal{W}$ denotes the Wishart distribution [12].

where μ_0 is a prior on μ , Σ_0 is a prior on Σ , and $\kappa, \nu \in \mathbb{R}$, $\kappa > 0$, $\nu > d - 1$. The likelihood of the M patches similar to the current patch (according to the L^2 distance) $\mathbf{z} = (z_1, \dots, z_M)$ is given by

$$p(\mathbf{z} | \mu, \Lambda) = \prod_{j=1}^M \frac{|\Lambda_j^*|^{-1/2}}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2} \sum_{j=1}^M z_j' \Lambda_j^* (z_j')^T\right), \quad (6)$$

with $z_j' = (z_j - D_j \mu)$ and $\Lambda_j^* = (D_j \Lambda^{-1} D_j^T + \Sigma_{N_j})^{-1}$. Then the MAP estimates $\hat{\mu}$ and $\hat{\Lambda}$ are found by maximizing the posterior probability

$$(\hat{\mu}, \hat{\Lambda}) = \arg \max_{\mu, \Lambda} p(\mu, \Lambda | \mathbf{z}) \quad (7)$$

$$= \arg \max_{\mu, \Lambda} p(\mathbf{z} | \mu, \Lambda) p(\mu, \Lambda). \quad (8)$$

Computing the partial derivatives of $p(\mu, \Lambda | \mathbf{z})$ with respect to μ and Σ and equating to zero we have (c.f. Appendix B in the supplementary material)

$$\hat{\mu} = \left(\kappa \mathbf{I} + \sum_{j=1}^M \Lambda^{-1} D_j \Lambda_j^* D_j \right)^{-1} \left(\sum_{j=1}^M \Lambda^{-1} D_j \Lambda_j^* z_j + \kappa \mu_0 \right) \quad (9)$$

$$\hat{\Lambda}^{-1} = H_1 H_2, \quad (10)$$

with

$$H_1 = \left((\nu - d) \mathbf{I} + \sum_{j=1}^M \Lambda^{-1} D_j \Lambda_j^* D_j \right)^{-1} \quad (11)$$

$$\begin{aligned} H_2 = & \left(\sum_{j=1}^M (\Lambda^{-1} D_j \Lambda_j^* z_j') (\Lambda^{-1} D_j \Lambda_j^* z_j')^T \right. \\ & \left. + \kappa (\mu - \mu_0) (\mu - \mu_0)^T + \nu \Sigma_0 \right). \end{aligned} \quad (12)$$

Since (9) depends on Λ and (10) depends on μ and Λ , those are not closed-forms for the estimators. Hence, we propose to use an iterative approach to compute the parameters, which is summarized in Algorithm 1. This estimation

algorithm results from the combination of two procedures. The outer loop follows from the classic EM estimation procedure for the mean and covariance (or precision) matrix. The inner one, which deals with the estimation of the precision matrix, converges if and only if the spectral norm of the precision matrix is less than one. This is not easy to prove analytically, given the expressions for H_1 and H_2 . In case this condition on the spectral norm of Λ holds, since EM and therefore the posteriors are guaranteed to converge to a local maxima, μ and Λ are ensured to converge to local maximizers. In practice, we observe that the algorithm converges after a single iteration of the outer loop with 3 to 4 iterations of the inner loop.

Algorithm 1: Computation of $\hat{\mu}$ and $\hat{\Lambda}$.

Input: $Z, D, \mu_0, \Sigma_0, \kappa, \nu$ (see details in Section 3.3)
Output: $\hat{\mu}, \hat{\Lambda}$

- 1 **Initialization:** Set $\Lambda = \Sigma_0^{-1}$
- 2 **for** $it = 1$ to $maxIts_0$ **do**
- 3 Compute $\hat{\mu}$ according to (9)
- 4 Set $\mu = \hat{\mu}$.
- 5 **for** $it = 1$ to $maxIts_1$ **do**
- 6 Compute $\hat{\Lambda}$ according to (10).
- 7 Set $\Lambda = \hat{\Lambda}$.
- 8 **end**
- 9 **end**

2.4. Summary of the proposed algorithm

The analysis previously presented leads to an iterative algorithm that implements the proposed approach. Two stages are alternated: the restoration step, where all patches are reconstructed, and the model estimation step, where the model parameters are updated (Figure 1). For the model estimation step, an *oracle* image is assumed to be available, i.e. the result of the previous iteration. In practice, the algorithm is found to converge after 3 to 4 iterations. The procedure is summarized in Algorithm 2.

3. Implementation details

3.1. Search for similar patches

The similar patches are all patches with L^2 distance to the current patch below a given threshold, which is given by a tolerance parameter times the distance to the closest neighbor. The patch comparison is performed in an oracle image (i.e. the result of the previous iteration), so all pixels are known. However, it may be useful to assign different confidence levels to the known pixels ($D_p^j = 1$) and to those originally missing and then restored ($D_p^j = 0$). For all the experimental results presented in Section 4, the distance between patches c_p and c_q in the oracle image C_{oracle}

Algorithm 2: Summary of the proposed algorithm.

Input: $Z, D, \mu_0, \Sigma_0, \kappa, \nu$ (see details in Section 3.3)
Output: $\hat{C}$

- 1 Decompose Z and D into overlapping patches.
- 2 **Initialization:** Compute first oracle image C_{oracle} (see details in Section 3.4)
- 3 **for** $it = 1$ to $maxIts_2$ **do**
- 4 **for** all patches not yet restored **do**
- 5 Find patches similar (L^2 distance) to the current z_i in C_{oracle} (see details in Section 3.1).
- 6 Compute μ_0 and Σ_0 from C_{oracle} (see details in Section 3.3).
- 7 Compute $\hat{\mu}$ and $\hat{\Sigma}$ following Algorithm 1.
- 8 Restore the similar patches using (4) (see details in Section 3.2).
- 9 **end**
- 10 Perform aggregation to restore the image.
- 11 Set $C_{oracle} = \hat{C}$.
- 12 **end**

is computed according to

$$d(p, q) = \frac{\sum_{j=1}^N (c_p^j - c_q^j)^2 \omega_{p,q}^j}{\sum_{j=1}^N \omega_{p,q}^j}, \quad (13)$$

with $\omega_{p,q}^j = 1$ if $D_p^j = D_q^j = 1$ and $\omega_{p,q}^j = 0.01$ otherwise. With this formulation, we are giving much higher priority to the known pixels compared to the unknown ones. Variations of these weights could be explored.

3.2. Computation of μ and Σ

The proposed method computes one Gaussian model per image patch according to Equations (9) and (10). In order to reduce the computational cost, we rely on the collaborative filtering idea previously introduced for patch-based denoising techniques [8, 6]. Based on the hypothesis that similar patches share the same model, instead of computing a different pair (μ, Σ) for each patch, we assign the same model to all patches in the set of similar patches. The restoration is thus performed for all similar patches according to the computed model.

3.3. Parameters setting

The four parameters of the Normal-Wishart distribution: κ, ν , the prior mean μ_0 and the prior covariance matrix Σ_0 , must be set in order to compute μ and Σ using (9) and (10).

Setting of κ and ν The computation of μ according to (9) combines the mean estimated from the similar patches and the prior mean μ_0 . The parameter κ is related to the degree

of confidence we have on the prior μ_0 . Hence, its value should be a trade-off between the confidence we have on the prior accuracy vs. the one we have on the information provided by the similar patches. A higher κ is needed if few similar patches are available or if a large part of the patch is unknown. Similarly, the ν parameter should be set to define a trade-off between the information provided by the similar patches and the prior Σ_0 . Despite these intuitive insights, setting these parameters is not a trivial task and should be the subject of further study. For the results presented in Section 4, the parameters are defined as follows. Because the importance we give to the priors μ_0 and Σ_0 is controlled by the relative importance of the value of κ and ν with respect to the diagonal values of the term $S = \sum_{j=1}^M \Lambda^{-1} D_j \Lambda_j^* D_j$ (see equations (9) and (11)), we set κ and ν proportional to the mean value of the diagonal entries of this matrix

$$\kappa, \nu = \frac{\alpha}{n} \sum_{i=1}^n S_i, \quad \alpha = \begin{cases} \alpha_L & \text{if } P \text{ and } M > \text{threshold} \\ \alpha_H & \text{otherwise,} \end{cases} \quad (14)$$

where P is the number of known pixels in the current patch and S_i is i -th diagonal element of matrix S .

Setting of μ_0 and Σ_0 Assuming an oracle image C_{oracle} is available (see details in Section 2.4), μ_0 and Σ_0 can be computed using the classical MLE estimators from a set of similar patches ($\tilde{c}_1, \dots, \tilde{c}_M$) taken from C_{oracle}

$$\mu_0 = \frac{1}{M} \sum_{j=1}^M \tilde{c}_j, \quad \Sigma_0 = \frac{1}{M-1} \sum_{j=1}^M (\tilde{c}_j - \mu_0)(\tilde{c}_j - \mu_0)^T. \quad (15)$$

This is the same approach followed by Lebrun et al. [8] to locally estimate the patch model parameters in the case of denoising. As previously stated, the method from [8] cannot be directly applied to zooming or interpolation due to the presence of missing pixels.

3.4. Initialization

A good initialization is crucial since we aim at solving a non-convex problem through an iterative procedure. Yu et al. [16] propose to initialize the PLE algorithm learning the K GMM covariance matrices from synthetic images of edges with different orientations as well as the DCT basis to represent isotropic patterns. As they state, in dictionary learning, the most prominent atoms represent local edges which are useful at representing and restoring contours. Hence, this initialization helps to correctly restore corrupted patches even in quite extreme cases. The first oracle of the proposed iterative approach is created by aggregating the estimations of all patches obtained through the initialization process proposed by Yu et al. [16]. Figure 1 illustrates the proposed initialization. A detailed description

of the initialization process is provided in the supplementary material.

4. Experiments

In this section we illustrate the ability of the proposed method to solve a number of diverse image inverse problems. Both synthetic and real data are used. The considered problems are: interpolation, denoising, and zooming. The reported values of peak signal-to-noise ratio ($PSNR = 20 \log_{10}(255/\sqrt{MSE})$) are the mean of 10 realizations for each experiment (variance is below 0.1 for interpolation and it is below 0.05 for denoising). The complete result images are included in the supplementary material.

4.1. Synthetic data

Interpolation Random masks with 20%, 50% and 70% of missing pixels are applied to the tested ground-truth images. The interpolation performance of the proposed method is compared to that of PLE [16], EPLL [18] and E-PLE [13] using a patch size of 8×8 for all methods. PLE parameters are set as indicated in [16] ($\sigma = 3$, $\varepsilon = 30$, $K = 19$). We used the EPLL code provided by the authors [17] with default parameters and the E-PLE code available in [13] with the parameters set as specified in this demo. The parameters for the proposed method are set to $\alpha_H = 1$, $\alpha_L = 0.5$ (α_H and α_L define the values for κ and ν , see Section 3.3). The PSNR results are shown in Table 1. Figure 2 shows some extracts of the obtained results, the PSNR values for the extracts and the corresponding difference images with respect to the ground-truth. The proposed method gives sharper results than the other considered methods. This is specially noticeable on the reconstruction of the texture of the fabric of Barbara's trousers shown in the first row of Figure 2 or on the strips that appear through the car's window shown in the second row of the same figure.

Denoising The following experiments are conducted in order to compare the denoising ability of the proposed method to that of the state-of-the-art denoising algorithm NLB [8]. The experiments are performed with images corrupted with additive Gaussian noise with variance $\sigma^2 = 10, 30, 50, 80$. The code provided by the authors [9] automatically sets the algorithm parameters from the input σ^2 and the patch size, in this case 8×8 . For this experiment, there are no unknown pixels to interpolate (the mask D is the identity matrix). The results of both methods are very similar if HBE is initialized with the output of the first step of NLB [8] (instead of using the initialization described in Section 2.4) and the parameters κ and ν are large enough. In that case, μ_0 and Σ_0 are prioritized in equations (9) and (10) and both algorithms are almost the same. That is what we observe in practice with $\alpha_H = \alpha_L = 100$, as exemplified in the results summarized in Table 1. The denoising

		PSNR (dB)											
		HBE	PLE	EPLL	E-PLE	HBE	PLE	EPLL	E-PLE	HBE	PLE	EPLL	E-PLE
% missing pixels		20%				50%				70%			
Interpolation	barbara	45.73	43.48	40.89	43.75	39.17	36.93	32.99	35.43	34.74	32.50	27.96	28.77
	boat	41.76	40.37	40.17	40.32	35.14	34.32	34.21	33.59	31.47	30.74	30.38	30.26
	traffic	35.73	35.53	35.71	35.10	30.28	30.12	30.19	28.86	27.41	27.12	27.13	26.64
		HBE	NLB	EPLL	HBE	NLB	EPLL	HBE	NLB	EPLL	HBE	NLB	EPLL
σ^2		10				30				50			
Denoising	barbara	41.39	41.20	40.56	38.57	38.26	37.32	37.32	36.94	35.84	36.15	35.73	34.51
	boat	40.24	39.99	39.47	36.83	36.76	36.34	35.58	35.46	35.13	34.51	34.33	34.12
	traffic	40.87	40.74	40.55	37.24	36.99	36.86	35.57	35.26	35.20	34.04	33.70	33.72

Table 1. Results of the interpolation and denoising tests described in Section 4.1. Patch size of 8×8 for all methods in all tests. Parameter setting for interpolation, HBE: $\alpha_H = 1$, $\alpha_L = 0.5$, PLE: $\sigma = 3$, $\varepsilon = 30$, $K = 19$ [16], EPLL: default parameters [17], E-PLE: parameters set as specified in [13]. Parameter setting for denoising, HBE: $\alpha_H = \alpha_L = 100$, NLB: code provided by the authors [9] automatically sets parameters from input σ^2 , EPLL: default parameters for the denoising example [17]

performance of HBE is degraded for small κ and ν values. The reason for this is that μ_0 and Σ_0 , as well as μ and Σ in NLB, are computed from an oracle image resulting from the first restoration step. This restoration includes not only the denoising of each patch, but also an aggregation step that highly improves the final result. Therefore, the contribution of the first term of (9) to the computation of $\hat{\mu}$ degrades the result compared to using μ_0 only (i.e. using a large κ).

Zooming In order to evaluate the zooming capacity of the proposed approach, ground-truth images are downsampled by a factor 2 (no anti-aliasing filter is used) and the zooming is compared to the ground-truth. The results are compared with PLE, EPLL, E-PLE and the Lanczos interpolation. Figure 3 shows extracts of the obtained results, the PSNR values for the extracts and the corresponding difference images with respect to the ground-truth. HBE yields a sharper reconstruction than the other methods.

4.2. Real data

Raw images are captured using a Canon 400D camera set to ISO 400 and exposure time 1/160 seconds. Because we treat raw data, pixels' noise have variable variance (according to Model (3)) and a camera calibration step is needed to compute the noise model parameters [2]. In order to evaluate the interpolation capacity of the proposed approach, we consider the pixels of the green channel only (i.e. 50% of the pixels in the RGGB Bayer pattern) and interpolate the missing values. We compare the results to those obtained using an adaptation of PLE to images degraded with noise with variable variance [1]. The results for the EPLL and E-PLE methods are not presented since these methods are not suited for this kind of noise. Figure 4 shows examples of the obtained results. As had already been observed in the synthetic data experiments, fine details and edges are better preserved. Compare for example the reconstruction of the

balcony edges and the wall structure in the first row of Figure 4, as well as the structure of the roof and the railing in the second row of the same image.

4.3. Discussion

In all the tested examples, the results obtained by HBE for every considered inverse problems outperform or are very close to those obtained by the other evaluated methods. Details are better reconstructed and the resulting images are more sharp both in the synthetic and real data examples. The improvement is more noticeable comparing the difference images (available for the synthetic tests only), which present less structure in the result obtained by HBE.

Even if the PLE method can be considered as semi-local (since it is applied in 128×128 regions [16]), we find that 19 classes are not enough to correctly represent certain image patches. This is mostly the case for patches that seldomly appear in the image, such as certain edges or particular textures that appear in a few patches. This is quite noticeable in the extract of barbara's trousers and in the interior of the car (Figure 2). The specific characteristics of these patches are buried in the PLE class update when combined with many other different patches. A local model estimation as the one performed by HBE correctly handles those cases. The performance difference is much more remarkable for the higher masking rates. In those cases, two phenomena take place. On the one hand, very few pixels are known thus making the model selection less robust. On the other hand, the model accuracy is critical since a much larger part of the patch is to be restored. The proposed method tackles the model selection problem by limiting the model estimation to similar patches found on a local search window. It has been widely observed in denoising techniques based on the self-similarity principle [4] that performance improves when restricting the patch search space to a local search window instead of using the whole image. This strategy, in addition




Figure 2. **Synthetic data. Interpolation with 70% of randomly missing pixels.** Left to right: (first row) Ground-truth (extract of barbara), result by HBE, PLE, EPLL, E-PLE. (second row) input image, difference with respect to the ground-truth of each of the corresponding results. (third and fourth row) Idem for an extract of the traffic image. See Table 1 for the PSNR results for the complete images. Please see the digital copy for better details reproduction.

to the hypothesis of self-similarity in that neighborhood, restricts the possible models robustifying the model estimation which is crucial for high masking rates. Furthermore, the local model estimation, previously proven successful at describing patches [8], gives a better reconstruction even when a very large part of the patch is missing.

EPLL uses more mixture components in its GMM model than PLE, where 200 components are learnt from 2×10^6 patches of natural images [18]. The results obtained by this approach, despite using a larger number of GMM com-

ponents, are not very good for the restoration of certain patches. As previously mentioned, Wang and Morel [15] claim that, in the case of denoising, it is better having fewer models that are updated with the image patches (as in PLE) than having a large number of fixed models (as in EPLL). In this work, we observe that the proposed approach outperforms EPLL, not only in denoising, but also in inpainting and zooming. However, it is here harder to tell if the improvement is due to the local model estimation performed from the similar patches or it is due to the different approach




Figure 3. **Synthetic data. Zooming $\times 2$. Left to right:** (first row) Ground-truth high resolution image (extract of lena). Result by HBE, PLE, EPLL, E-PLE, lanczos interpolation. (second row) Input low-resolution image, difference with respect to the ground-truth of each of the corresponding results. Please see the digital copy for better details reproduction.



Figure 4. **Real data. Zooming $\times 2$.** Interpolation of the green channel of a raw image (RGGB). **Left to right:** Input low-resolution image, result by HBE, PLE adapted to noise with variable variance [1], bicubic and lanczos interpolation.

that both methods follow for restoration.

5. Conclusions

In this work we have presented a novel image restoration framework. It has the benefits of local patch characterization proven by the NLB denoising methods, but manages to

extend its use to general restoration problems such as zooming, inpainting and interpolation, by combining a local estimation with a Bayesian restoration based on hyperpriors. We presented various experiments that confirm the soundness of the proposed strategy based on hyperpriors. In these experiments, using both synthetic and real data, we show that for a wide range of image restoration problems, HBE outperforms several state-of-the-art restoration methods.

References

- [1] C. Aguerrebere, A. Almansa, J. Delon, Y. Gousseau, and P. Muse. Single shot high dynamic range imaging using piecewise linear estimators. In *Proceedings of IEEE International Conference on Computational Photography (ICCP)*, 2014. 6, 8
- [2] C. Aguerrebere, J. Delon, Y. Gousseau, and P. Musé. Study of the digital camera acquisition process and statistical modeling of the sensor raw data. *Technical report <hal-00733538v1>*. 1, 2, 6
- [3] A. Buades, B. Coll, and J. M. Morel. A Review of Image Denoising Algorithms, with a New One. *SIAM Multiscale Modeling & Simulation*, 4(2):490–530, 2005. 1
- [4] A. Buades, B. Coll, and J.-M. Morel. Non-Local Means Denoising. *Image Processing On Line*, 1, 2011. 6
- [5] P. Chatterjee and P. Milanfar. Patch-Based Near-Optimal Image Denoising. *IEEE Transactions on Image Processing*, 21(4):1635–1649, 2012. 1
- [6] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8), 2007. 4
- [7] A. A. Efros and T. K. Leung. Texture Synthesis by Non-Parametric Sampling. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, pages 1033–, 1999. 1
- [8] M. Lebrun, A. Buades, and J. Morel. A nonlocal bayesian image denoising algorithm. *SIAM Journal on Imaging Sciences*, 6(3):1665–1688, 2013. 1, 2, 3, 4, 5, 7
- [9] M. Lebrun, A. Buades, and J.-M. Morel. Implementation of the “Non-Local Bayes” (NL-Bayes) Image Denoising Algorithm. *Image Processing On Line*, 3:1–42, 2013. 5, 6
- [10] M. Lebrun, M. Colom, A. Buades, and J. Morel. Secrets of image denoising cuisine. *Acta Numerica*, 21(1):475–576, 2012. 1
- [11] S. Lyu and E. Simoncelli. Modeling Multiscale Subbands of Photographic Images with Fields of Gaussian Scale Mixtures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(4):693–706, 2009. 1
- [12] H. Raiffa and R. Schlaifer. *Applied statistical decision theory [by] Howard Raiffa and Robert Schlaifer*. Division of Research, Graduate School of Business Administration, Harvard University Boston, 1961. 3
- [13] Y.-Q. Wang. E-PLE: an Algorithm for Image Inpainting. *Image Processing On Line*, 2013:271–285, 2013. 2, 5, 6
- [14] Y.-Q. Wang. The Implementation of SURE Guided Piecewise Linear Image Denoising. *Image Processing On Line*, 3:43–67, 2013. 1
- [15] Y.-Q. Wang and J.-M. Morel. SURE Guided Gaussian Mixture Image Denoising. *SIAM Journal on Imaging Sciences*, 6(2):999–1034, 2013. 1, 2, 7
- [16] G. Yu, G. Sapiro, and S. Mallat. Solving inverse problems with piecewise linear estimators: From gaussian mixture models to structured sparsity. *IEEE Transactions on Image Processing*, 21(5):2481–2499, 2012. 1, 2, 5, 6
- [17] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. <http://people.csail.mit.edu/danielzoran/epllcode.zip>. Accessed: 29/09/2014. 5, 6
- [18] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, pages 479–486, 2011. 1, 2, 5, 7