Presented at ICIP 2006, vol I, pp. 625-628,October 2006, Atlanta.

DEBLURRING-BY-DENOISING USING SPATIALLY ADAPTIVE GAUSSIAN SCALE MIXTURES IN OVERCOMPLETE PYRAMIDS

Jose A. Guerrero-Colon * *and Javier Portilla* †

Visual Information Processing Group
Dept. of Computer Science and Artificial Intelligence
Universidad de Granada
{jaguerrero, javier}@decsai.ugr.es

In a previous work, we presented an extension of the original Bayes Least Squares - Gaussian Scale Mixtures (BLS-GSM) denoising algorithm that also compensated the blur. However, that method had some problems: a) it could not compensate for some blurring kernels; b) its performance depended critically on having an accurate estimation of the original power spectral density (PSD); and c) it could not be easily adapted to a spatially variant description of the image statistics. In this work we propose a two-step restoration method that overcomes these problems by first performing a global blur image compensation, and then applying a spatially adaptive local denoising, in an overcomplete pyramid. Our method is efficient, robust and non-iterative. We demonstrate through simulations that it provides state-of-the-art performance.

Index Terms— Image restoration, wavelet transforms

1. INTRODUCTION

Image restoration is a classical problem which most often aims to estimate an image given a linearly filtered, noisy version of it. This is a difficult ill-conditioned inverse problem, even if, as in this work, the degradation parameters are assumed known. The Bayesian estimation, by making an explicit formulation of our prior knowledge about natural images, provides a powerful tool for this task. Typically, restoration methods are iterative and global, e.g. [1, 2, 3]. Among global methods, some authors (e.g. [4, 5, 6, 7, 8]) have approached the deconvolution problem by decomposing it in two steps (see Fig. 1): 1) a global blur compensation, and 2) a pure denoising step. The main advantage of this decomposition is that the deconvolution problem is so converted into an easier non-white noise removal problem, at the likely price of losing global optimality. [1] Given the increasingly high performance of current denoising methods, this approach has become an appealing practical alternative.

*JAGC is supported by 2245-1065 grant contract, funded by Ilía Sistemas S.L. and Ministerio de Defensa (Spain).

†JP is supported by TIC2003-01504 grant, funded by the Ministerio de Educación y Ciencia (Spain) and the "Ramón y Cajal" program.

[1] Note that even if we model noise as white, it will no longer be white after the blur compensation.

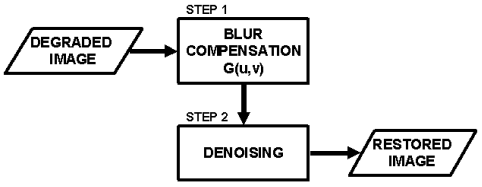

Fig. 1: Restoration Scheme.

A less common path (see e.g. [9, 10]), has been to approach the restoration problem by using local models. A local description of the signal may simplify the formulation and the resolution of the problem, but it also imposes some limitations: how can we compensate for a large blurring kernel using a small neighborhood for the local description? As we will see in Section 2, multi-resolution does not give a complete solution. In [10] we used a local GSM [11] model which provides high performance using an efficient non-iterative method. However this approach has problems inherent to its local nature, and it requires the estimation of the original spectral density from the degraded observation.

In this work we propose a two-step deblurring method based on a recent spatially adaptive local GSM model [12]. By choosing properly the global pre-filtering, our method achieves a close to optimal trade-off between (global) blur compensation and (local) noise removal, improving in performance and robustness with respect to previous methods.

2. DEBLURRING USING BLS-GSM

Our observation is modelled as follows:

$$\mathbf{y}_0 = \mathbf{H}\mathbf{x} + \mathbf{w}_0, \tag{1}$$

where $\mathbf{H}$ is a block-circulant matrix implementing a convolution with h (the Point Spread Function, PSF, typically a low-pass filter) and the vectorized (following a lexicography order) original signal $\mathbf{x}$. The vector $\mathbf{w}_0$ is Gaussian noise of known PSD. We search for a plausible estimate of $\mathbf{x}$ given the observation, assuming that both the noise PSD ($P_{\mathbf{w}_0}$) and h are known.

2.1. One step Deblurring

The power of the BLS-GSM approach for denoising has been widely proved [13, 12]. The same approach has been used for image restoration in [10] with good results. The strength of that restoration method relies on using a local model for spatial neighbors within subbands of overcomplete pyramids. The pyramidal representation, together with the local filtering, results in a very efficient method. However, using a local model imposes also some limitations. Given a small neighborhood, we can only compensate for smoothly varying (in frequency) filters. A fixed neighborhood size for all the scales gives raise to larger effective neighborhoods as we go up in the pyramid. So, thanks to multi-resolution, we can use (and compensate for) smooth filters in frequency at low scales, less smooth at medium scales and peaky at high scales (low frequencies). Such behavior is suitable for many low-pass kernels, like Gaussian or Laplacian. However, there are some practically important kernels (e.g., uniform) whose behavior in frequency is the other way around, i.e., smooth at low frequencies and oscillating at medium/high frequencies. For those cases, the local multi-scale blur compensation fails. Another issue is that [10] requires an accurate estimation of the original PSD ($P_{\mathbf{x}}$). This is a serious problem, especially if we want to make the model spatially adaptive by using blocks, as we did in [12] from [13].

2.2. Two-step Deblurring

To preserve the strength of the referred local model while overcoming its problems, we use the scheme depicted in Fig. 1, followed by other authors [4, 6, 7, 8]. Our steps are: 1) global blur compensation and, 2) local denoising of the compensated observation. We have experienced that this approach increases the robustness because it reduces the dependency on the estimated $P_{\mathbf{x}}$, and it allows us to take advantage of a state-of-the-art spatially variant local GSM denoising [12].

In order to get an LS-optimal result, we should properly couple both steps. From Eq. 1 we can express the observation model in the frequency domain as

$$Y_0(u,v) = H(u,v)X(u,v) + W_0(u,v) \tag{2}$$

where[2] H is the Fourier transform of kernel h. Let G be the filter applied in step 1. We define E^2_{step-2}, an estimate of the Mean Square Error (MSE) of the denoising block output w.r.t. the original, as a function of its input, its equivalent PSF, and its noise PSD. The optimal G is, thus,

$$G_{opt} = \arg\min_{G} E^2_{step-2}(GY_0, GH, |G|^2 P_{\mathbf{w}_0}). \tag{3}$$

At this time we are still working on an MSE model of the BLS-GSM denoising algorithm. Meanwhile, we propose the following simplified approach:

[2]For notational simplicity we will drop the indices (u,v).

Step 1: *Global blur compensation*

This step provides a blur-compensated image by applying a global filtering to the observed image. Instead of estimating the optimal filter as a whole (Eq. 3), we reduce the degrees of freedom to only one, called the regularization parameter, α. We use, as in other works [5, 6] a generalized Wiener filter:

$$G = \frac{H^*}{|H|^2 + \alpha \cdot \left(\frac{P_{\mathbf{w}_0}}{P_{\mathbf{x}}}\right)}. \tag{4}$$

The α parameter controls the noise/blur suppression. Low values result in very noisy intermediate images ($\alpha = 0$ is just an inverse filter) whereas high values smooth out the high frequencies. For $\alpha = 1$ the compensation is the Wiener filtering, which provides too smooth global restoration results. Therefore, we will consider α values lying in the $(0, 1)$ interval.

Step 2: *Local denoising*

After step-1, the observation yields, $Y_1 = H_r X + W_1$, where $H_r = GH$ and $W_1 = GW_0$. We could approach the problem as another deblurring step, but in practice the residual blur H_r is small and can be neglected. For the second step we have used our spatially adaptive BLS-GSM denoising algorithm [12], which is robust, non-iterative and provides very high performance. The only difference compared to [13] is that the latter assumes a fixed signal covariance for all GSM neighborhoods in a subband, and the former uses a signal covariance that changes at different regions (blocks). The observed neighborhood centered in position (n, m) at subband j is $\mathbf{y}_1^{j,(n,m)} = \mathbf{x}^{j,(n,m)} + \mathbf{w}_1^{j,(n,m)} = \sqrt{z}\mathbf{u}^{j,(n,m)} + \mathbf{w}_1^{j,(n,m)}$ where $\mathbf{u}^{j,(n,m)}$ and $\mathbf{w}_1^{j,(n,m)}$ are zero mean Gaussian vectors with covariance matrices $\mathbf{C}_{\mathbf{u}}^{j,(n,m)}$ and $\mathbf{C}_{\mathbf{w}_1}^{j,(n,m)}$, and $\sqrt{z}$ is the hidden multiplier controlling the variance of the local signal $\mathbf{x}^{j,(n,m)}$. The BLS-GSM method consists of estimating the original central coefficient from the observed neighborhood through:

$$\mathbb{E}\{x_c|\mathbf{y}_1\} \quad = \quad \int_0^\infty p(z|\mathbf{y}_1)\, \mathbb{E}\{x_c|\mathbf{y}_1, z\}\, dz \tag{5}$$

The expected value inside the integral is a local Wiener estimate, and the posterior $p(z|\mathbf{y}_1)$ is computed in terms of $p(\mathbf{y}_1|z)$ (Gaussian) and the Jeffrey's prior on z [13, 12].

2.3. Pre-filtering parameter estimation

For the $P_{\mathbf{x}}$ we have used a simple model that falls as k/f^2 in the frequency domain [14]. We have made the model consistent with the estimated original variance ($\hat{\sigma}^2_{\mathbf{x}} = \sigma^2_{\mathbf{y}_0} - \sigma^2_{\mathbf{w}_0}$) by setting $k = \hat{\sigma}^2_{\mathbf{x}} / \int_{f>0}(|H(f)|^2/f)\, df$. Despite its simplicity, this model is effective, and, unlike in [10], we have experienced that the global performance does not depend critically on an accurate $P_{\mathbf{x}}$ estimation.

Some methods [4, 6] use a model to choose the regularization parameter of the blur compensation. Lacking a proper

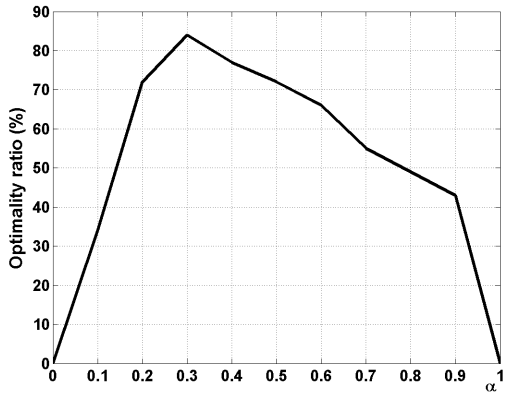

Fig. 2: Optimality ratio for each α value.

model, we have done it empirically, training with a set of five images (*House*, *Cameraman*, *Barbara*, *Boat*, and *Lena*) and Gaussian PSFs and noise of 0.4 to 3.2 and 1 to 16 standard deviations, respectively (in logarithmic steps). For each blur/noise combination we have computed the optimal α using the golden section algorithm and an interval around it with the α values providing a decrease in performance less than 0.05 dBs w.r.t. the optimum. Then we have chosen the common α value that provides the highest proportion of results within the intervals. Fig. 2 shows the optimality ratio for each alpha value. We have obtained $\alpha_{opt} = 0.3$, achieving roughly 84% of optimality. This high value means that in most cases we obtain results very close to optimal. We have also tested that this result is fairly independent on the particular images chosen for the training, so there is no significant *overfitting*.

3. RESULTS AND DISCUSSION

For all the experiments, we implemented step 1 in the Fourier domain, using $\alpha_{opt} = 0.3$. For step 2 we have used two representations: Trapezoidal Haar Pyramid (THP) and Full Steerable Pyramid (FSP). FSP is an extension of Simoncelli's original steerable pyramid [13] and THP is an overcomplete version of the Haar wavelet [12]. We have used 3 and 8 orientations for THP and FSP, respectively, and 4 scales. The representation choice depends on the relative amount of texture: texture-rich images are better denoised using FSP, whereas THP is superior for poorly textured images [12]. We have used a 3×3 GSM neighborhood and a 32×32 local block.

In Table 1 we show, in terms of averaged Increment in Signal-to-Noise Ration (ISNR), the result of applying our method to the training data set (see Section 2.3). In Table 2 we show a comparison to several state-of-the-art methods with six different experiments over *House*, *Cameraman* (using THP) and *Barbara* (using FSP). We compare with the previous BLS-GSM restoration method [10], and with whom, to the best of our knowledge, is the best-performing method [2][3]. The comparison can be extended to other methods [15, 3, 4, 6, 8, 16],

[3]We are very grateful to Prof. Figueiredo for providing us with his results.

		NOISE				
	σ	1	2	4	8	16
B	0.4	2.86	1.86	2.81	5.03	7.63
L	0.8	6.65	4.84	3.73	3.92	5.85
U	1.6	3.87	3.18	2.72	3.01	4.73
R	3.2	3.21	2.85	2.56	2.66	3.87

Table 1. Averaged results as Increment of Signal-to-Noise Ratio (ISNR), in dB., of the training data set (see sec. 2.3).

Blur	*PSF 1*		*PSF 2*	*PSF 3*	*PSF 4*	*PSF 5*
$\sigma_w \Rightarrow$	$\sqrt{2}$	$\sqrt{8}$	$\sqrt{0.308}$	7	2	8
Method	*HOUSE*					
[2]	8.47	6.63	**10.71**	4.22	**4.49**	4.76
[10]	8.46	6.93	-0.44	**4.37**	4.34	5.98
Ours	**8.64**	**7.03**	9.04	4.30	4.11	**6.02**
Method	*CAMERAMAN*					
[2]	**7.46**	5.24	**8.16**	2.84	3.18	3.65
[15]	6.93	4.88	7.59	**2.94**	-/-	-/-
[3]	7.40	5.15	8.10	2.85	-/-	-/-
[10]	6.84	5.29	-1.61	2.56	2.83	3.81
Ours	7.45	**5.55**	7.33	2.73	**3.25**	**4.19**
Method	*BARBARA*					
[2]	3.76	1.99	3.98	0.9	0.92	2.55
[10]	5.70	3.28	-0.27	1.44	0.95	4.91
Ours	**6.85**	**3.80**	**5.07**	**1.94**	**1.36**	**5.27**

Table 2. Performance comparison in terms of Increment of Signal-to-Noise ratio (ISNR), in dB. First row shows the blurring kernel used (see text for details), and the second denotes noise standard deviation. Best results are highlighted.

the best results among them have been included in Table 2. In the first two columns we have replicated the experiments in [8]. The kernel used (PSF1) is $h_{i,j} = (1 + i^2 + j^2)^{-1}$, for $i, j = -7...7$. Next column reproduces the experiment in [4, 6], that is a 9×9 uniform kernel (PSF2). Fourth column uses a 5×5 separable kernel (PSF3) with coefficients $[1, 4, 6, 4, 1]/16$ as in [16]. The last two columns correspond to Gaussian kernels with 1.6 and 0.4 standard deviation (PSF4 and PSF5, respectively). The noise standard deviations are shown in the second row of the table. These results show that [10] performs best in a 5.6% of the cases, [2] is best in a 22.2% and the proposed method is best in a 66.7% of the experiments. Fig. 3 shows a visual comparison of *House* and *Barbara* with the experiments of columns 1 and 3 from Table 2, respectively. As we can see, the roof line in our result is sharper than in the other two. *Barbara* results shows the fall in performance of [10] due to its inability to compensate for the blur of a uniform kernel (see Section 2.1). As visible in the tablecloth, our result recovers much more information than the competitors. The improvement is remarkable, both in visual and PSNR terms.

About computational cost, our method takes about 30 sec. with THP and 70 sec. with FSP in a 2.0GHz PC using Matlab© on a 256×256 image, whereas our most direct competitor [2] requires on average a similar time under the same conditions,

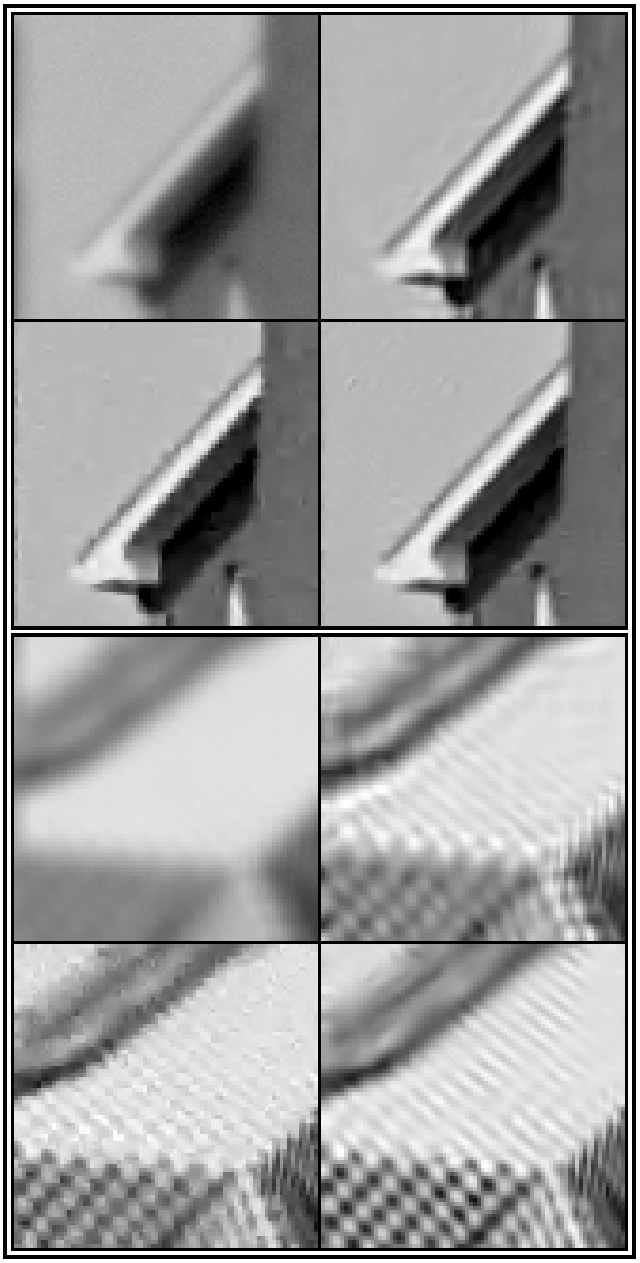

Fig. 3. Visual comparison results on *House* and *Barbara* images (cropped to 80×80). From left to right and top to bottom: Degraded (*House*: PSF1 & $\sigma_w = \sqrt{2}$; *Barbara*: PSF2 & $\sigma_w = \sqrt{0.308}$); BLS-GSM [10] with FSP; [2]; Our method (House with THP, *Barbara* with FSP). In the same order, Peak Signal-to-Noise ratio in dB. (*House* / *Barbara*): 25.62 / 22.49; 34.08 / 22.22; 34.09 / 26.47; 34.25 / 27.56.

but, unlike our method, its amount of computation depends on the degradation parameters.

Summarizing our method is robust and efficient, achieving state-of-the-art performance over a wide range of types and amounts of degradation. We still expect a significant increase in performance when including an LS-optimal prefiltering, based on an MSE model for our denoising method, instead of the empirical approach used here.

4. REFERENCES

[1] R. Molina, J. Mateos, A.K. Katsaggelos, and M. Vega, "Bayesian multichannel image restoration using compound gauss-markov random fields," *IEEE Trans. Image Proc.*, vol. 12, no. 12, pp. 1642– 1654, December 2003.

[2] M. Figueiredo and R. Nowak, "A bound optimization approach to wavelet-based image deconvolution," in *IEEE Int'l Conf on Image Proc*, 2005, vol. 2, pp. 782 – 785.

[3] J. Bioucas-Dias, "Bayesian wavelet-based image deconvolution: a gem algorithm exploiting a class of heavy-tailed priors," *IEEE Trans. Image Proc.*, vol. 4, pp. 937–951, April 2006.

[4] M. R. Banham and A. K. Katsaggelos, "Spatially adaptive wavelet-based multiscale image restoration," *IEEE Trans. Image Proc.*, vol. 5, pp. 619–634, Apr. 1996.

[5] R. Neelamani, H. Choi, and R. G. Baraniuk, "Wavelet-domain regularized deconvolution for ill-conditioned systems," in *IEEE Int'l Conf on Image Proc*, Kobe, Japan, Oct. 1999, vol. 1, pp. 204–208.

[6] R. Neelamani, H. Choi, and R. G. Baraniuk, "ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems," *IEEE Trans. Signal Proc.*, vol. 52, no. 2, pp. 418–433, Feb. 2004.

[7] J Kalifa and S Mallat, "Mini-max restoration and deconvolution," in *Bayesian inference in wavelet based methods*. Springer, 1999.

[8] A. Jalobeanu, N. Kingsbury, and J. Zerubia, "Image deconvolution using hidden markov tree modeling of complex wavelet packets," in *IEEE Int'l Conf on Image Proc*, 2001, vol. 1, pp. 201–204.

[9] AW Stevenson TE Gureyev, YI Nesterets and SW Wilkins, "A method for local deconvolution," *Applied Optics*, vol. 42, no. 32, pp. 6488–6494, November 2003.

[10] J Portilla and E P Simoncelli, "Image restoration using Gaussian scale mixtures in the wavelet domain," in *Proc IEEE Int'l Conf on Image Proc*, Barcelona, Spain, September 2003, vol. 2, pp. 965–968, IEEE Computer Society.

[11] D Andrews and C Mallows, "Scale mixtures of normal distributions," *J. Royal Stat. Soc.*, vol. 36, pp. 99–102, 1974.

[12] J.A. Guerrero-Colon and J. Portilla, "Two-level adaptive denoising using Gaussian scale mixtures in overcomplete oriented pyramids," in *IEEE Int'l Conf on Image Proc*. Genoa, Italy, September 2005, vol. I, pp. 105–108.

[13] J Portilla, V Strela, M Wainwright, and E P Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Trans. Image Proc.*, vol. 12, pp. 1338–1351, November 2003.

[14] D L Ruderman, "The statistics of natural images," *Network: Computation in Neural Systems*, vol. 5, pp. 517–548, 1996.

[15] M. Figueiredo and R. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. Image Proc.*, vol. 12, no. 8, pp. 906–916, Aug. 2003.

[16] J. Liu and P. Moulin, "Complexity-regularized image restoration," in *IEEE Int'l Conf on Image Proc*, 1998, vol. 1, pp. 555–559.