

Image restoration using sparse approximations of spatially varying blur operators in the wavelet domain

Paul Escande¹, Pierre Weiss¹, and François Malgouyres²

¹ ITAV-USR3505, Université de Toulouse, CNRS, Toulouse, France.

² IMT-UMR5219, Université de Toulouse, CNRS, Toulouse, France.

E-mail: paul.escande@gmail.com, pierre.armand.weiss@gmail.com,
francois.malgouyres@math.univ-toulouse.fr

Abstract. Restoration of images degraded by spatially varying blurs is an issue of increasing importance in the context of photography, satellite or microscopy imaging. One of the main difficulty to solve this problem comes from the huge dimensions of the blur matrix. It prevents the use of naive approaches for performing matrix-vector multiplications. In this paper, we propose to approximate the blur operator by a matrix sparse in the wavelet domain. We justify this approach from a mathematical point of view and investigate the approximation quality numerically. We finish by showing that the sparsity pattern of the matrix can be pre-defined, which is central in tasks such as blind deconvolution.

1. Introduction

This problem of image restoration in the presence of spatially varying blurs drew a considerable attention in the context of satellite imaging with the Hubble space telescope [1]. It is now becoming increasingly important in various imaging modalities such as 3D fluorescence microscopy [2]. In this paper, we aim at providing new computational tools to tackle this problem. We consider a blurring operator H in $\Omega \subset \mathbb{R}^d$, defined for any $u \in \mathbb{L}^2(\Omega)$ as the following integral operator,

$$\forall x \in \Omega, \quad Hu(x) = \int_{y \in \Omega} K(x,y)u(y)dy, \quad (1)$$

where $\Omega \subset \mathbb{R}^d$ is the image domain. The function K is a spatially varying kernel defining the Point Spread Function (PSF) at each location $x \in \Omega$.

The naive computation of a product Hu is simply obtained by direct discretization of (1) and costs $\mathcal{O}(N^{2d})$ arithmetic operations for an image of dimension d with N pixels in each dimension. It is unsuitable for large scale problems encountered in imaging. Previous works attempted to reduce this computational cost by approximating H using either tensor products or piecewise convolutions [3, 1, 4]. Unfortunately, these strategies do not have a property that is highly desirable in numerical analysis: the ability to approximate the original operator with an arbitrary precision. In practice, we observed that they are often too coarse for complex situations encountered in microscopy imaging.

In a recent work [5], we proposed to approximate H by an operator diagonal in the wavelet domain. However, this approximation was shown to be too coarse for some practical applications. Our main contribution in this paper consists of showing that the diagonal approximation can be replaced by a *sparse* approximation and lead to arbitrarily close approximations. More precisely, we show that $\|H - \Psi \Theta \Psi^*\|$, can be made arbitrarily small if Ψ^* denotes a forward wavelet transform and S is a sparse matrix that

contains only $O(N^d)$ non zero coefficients instead of N^{2d} . As a practical consequence, matrix-vector products using the proposed approach costs $O(d \log(N)N^d)$ arithmetic operations, which is doable even for very large images. We also show that the sparsity pattern of S can be known a priori, which is a highly desirable property in the case of blind deconvolution, where the operator should be inferred from the blurred images.

In section 2, we recall the theoretical results motivating this approximation. Section 3.1, contains preliminary restoration results in order to illustrate theoretical results and to identify a number of coefficients offering a compromise between the storage costs and the restoration quality. In this section, the kept coefficients are obtained by simple thresholding and can have arbitrary coordinates. In section 3.2, we finally show that the sparsity pattern of the matrix Θ can be pre-defined and that this approach leads to almost equivalent results to a thresholding.

2. Theoretical motivations

Sparse approximations of integral operators have been theoretically analyzed in [6, 7]. Surprisingly this approach was never applied to the approximation of spatially varying blur operators. The closest application found in the literature deals with foveation [8], but it is not adapted to a large class of kernels.

We define an orthonormal wavelet basis of $\mathbb{L}^2(\mathbb{R})$ as the family of functions:

$$\{\Psi_{j,m}\}_{j,m \in \mathbb{Z}}.$$

Each $\Psi_{j,m}$ is a dilated and translated of the mother wavelet Ψ ,

$$\Psi_{j,m}(t) = \sqrt{2^{-j}}\Psi(2^{-j}t - m).$$

Let us recall a typical result that motivate the proposed approach. We stick to the one-dimensional case for the ease of exposition. Since H is a linear operator in a Hilbert space, it can be written as $H = \Psi\Theta\Psi^*$, where $\Theta : l^2 \rightarrow l^2$ is the matrix representation of the blur operator in the wavelet domain. Matrix Θ is characterized by the coefficients:

$$\theta_{j,m,k,n} := (\langle H\Psi_{j,m}, \Psi_{k,n} \rangle), \quad \forall j, m, k, n \in \mathbb{Z}.$$

Theorem 1 (Decay of $\theta_{j,m,k,n}$ – [6]). *Supposing that:*

- the wavelets are compactly supported with $\text{supp}(\Psi_{j,m}) = I_{j,m}$,
- the wavelets have M vanishing moments,
- the operator H belongs to the class of Calderon-Zygmund operators, meaning that K satisfies:

$$\exists C_M > 0, \text{ such that } \forall x, y \in \Omega, \quad \begin{cases} |K(x, y)| \leq \frac{1}{|x - y|}, \\ |\partial_x^M K(x, y)| + |\partial_y^M K(x, y)| \leq \frac{C_M}{|x - y|^{1+M}}, \quad M \leq 1 \end{cases} \quad (2)$$

then the coefficients $\theta_{j,m,k,n}$ satisfy the following estimate:

$$|\theta_{j,m,k,n}| \leq C_M \min \left\{ \frac{|I_{j,m}|}{|I_{k,n}|}, \frac{|I_{k,n}|}{|I_{j,m}|} \right\}^{\frac{1}{2}} \left(\frac{\min \{|I_{j,m}|, |I_{k,n}|\}}{\text{dist}(I_{j,m}, I_{k,n})} \right)^{M+1}, \quad (3)$$

with C_M a constant depending on M and the wavelets, $\text{dist}(I_{j,m}, I_{k,n})$ denotes the distance between the two supports of the wavelets $\Psi_{j,m}$ and $\Psi_{k,n}$.

Moreover, if the kernel K is compactly supported, then for sufficiently large $|m - n|$,

$$\theta_{j,m,k,n} = 0. \quad (4)$$

This results ensures that the coefficients away from the diagonal fastly decay to zero. A more precise analysis shows that $O(dN^d \log(N))$ non zero coefficients allow to approximate H with an arbitrary precision [6]. Figure 1 shows the matrix Θ associated to a typical spatially varying blur operator in log scale. It is readily seen that the decrease of the coefficients away from the diagonal is extremely fast.

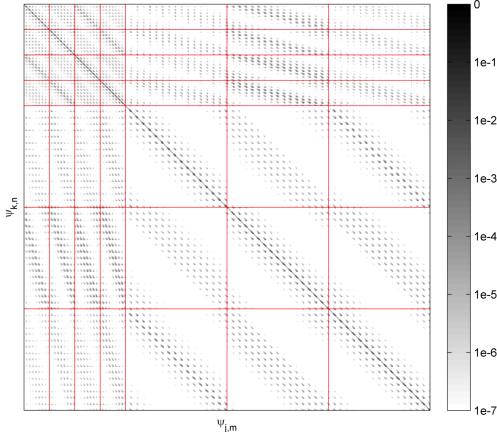


Figure 1. An example of a Θ matrix in a \log_{10} scale for a 256×256 images. The matrix contains about 4.3 billion coefficients.

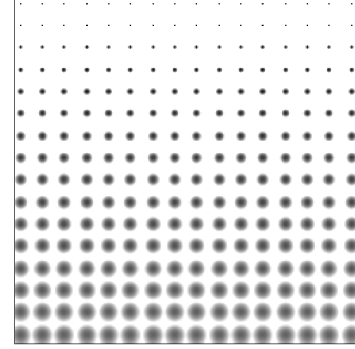


Figure 2. The point spread function of the blur used to generate the matrix on the left. The kernel is a Gaussian with equal variances increasing in the vertical direction.

In the next section, we reconstruct degraded images from the knowledge of various thresholded Θ matrices.

3. Restoration results

3.1. Restoration using thresholded matrices

In a first experiment, the Θ matrix is thresholded in order to keep a number of coefficients being a multiple of N^d (the number of pixels of the image). That is, we define the thresholded matrix Θ_{T_k} by zeroing all but the $T_k = kN^d$ largest coefficients of Θ . This experiment allows the identification of an admissible number of coefficients offering a proper compromise between the storage costs and the image restoration quality. To achieve this experiment, a full Θ matrix is numerically computed and then successively thresholded for various $k \in \mathbb{N}^*$. Thereafter, the image is restored assuming the following classical degradation model:

$$v = Hu + \eta, \quad \eta \sim \mathcal{N}(0, \sigma^2 I_{n^d}),$$

where v is the degraded image observed (see Figure 3c), u is the image to restore (see Figure 3a), and H is a blurring operator corresponding to the PSF displayed in Figure 3d. A standard TV-L2 optimization problem is solved to restore the image u :

$$\text{Find } u^* \in \underset{u \in \mathbb{R}^{n^d}, \|\Psi \Theta_{T_k} \Psi^* u - v\|_2 \leq \sigma^2 n}{\text{arg min}} \quad TV(u), \quad (5)$$

where $TV(u)$ is the isotropic total variation of u .

Figure 4 displays restored images for various values of k and the sparsity pattern of Θ_{T_k} . As expected, considering more coefficients improves the deblurring results. However, in the presence of noise, the two images restored using the full matrix and the thresholded matrix with $k = 20$, provide very similar results. Therefore, it is possible to store only a small amount of coefficients and obtain results close to the ones obtained having a full (unrealistic) knowledge of the blur matrix. Furthermore, these results are obtained drastically faster, see the speed-up Figure 4.

Remark: it is near impossible to solve the deblurring problem with an exact operator due to the numerical burden of computing matrix-vector products in large dimensions. We thus chose a spatially varying blur that is sparse in the wavelet domain in order to be able to make comparisons between an exact computation and an approximated one. Speed up times might thus be much higher for other kinds of blurs.

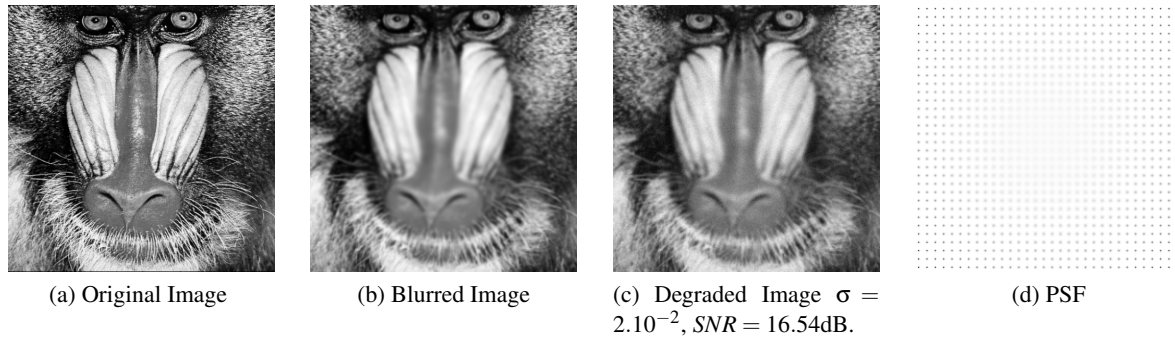


Figure 3. Images involved in the deblurring problem

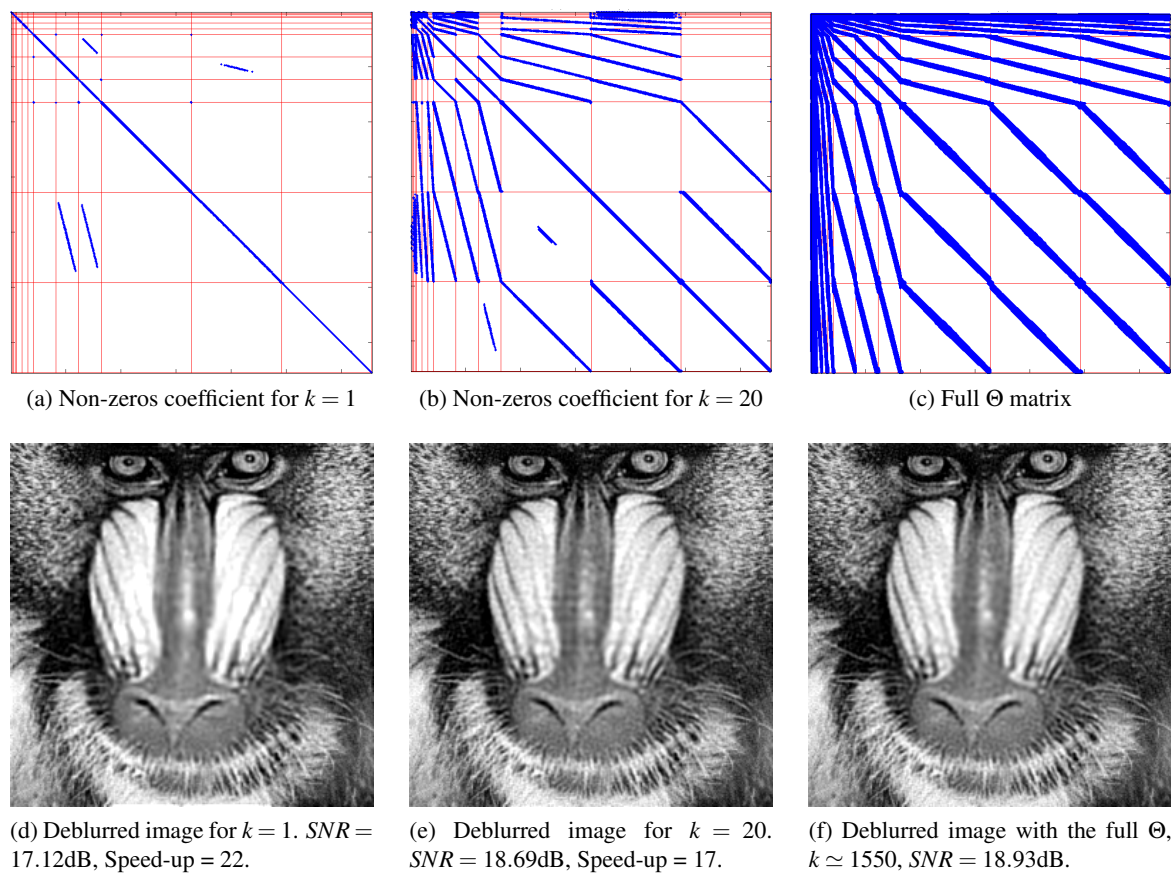


Figure 4. Deblurred images using thresholded matrices

3.2. Restoration using pre-defined sparsity patterns

We are currently developing strategies to solve blind deconvolution problems using the proposed approximation. In this setting, it is not possible to threshold the exact (unknown) matrix to obtain a sparsity pattern. The knowledge of such a structure is primordial to approximate K using prior information. Our aim in this section is to show that sparsity patterns can be pre-defined, allowing convincing restoration results.

Our main observation is that for a given wavelet $\psi_{j,m}$, only its neighbours in the same scale, the finer scale and the coarser one lead to significant correlation coefficients $\langle H\psi_{j,m}, \psi_{k,n} \rangle$. This concept of neighbours is also supported by Theorem 1. Thus, sparsity patterns can be defined using notions of intra and inter-scale neighbourhoods.

For each node of the wavelet decomposition quadtree, a neighbourhood is defined. It describes which correlation coefficients shall be preserved to generate the sparsity pattern. Figure 5 illustrates this idea. In this example, we consider a wavelet transform of depth 2. A neighbourhood \mathcal{N}_i is associated to each sub-band. For instance \mathcal{N}_1 can be represented as follows:

$$\mathcal{N}_1 = \begin{pmatrix} 2 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ l & l & l & l & l & h & v & d \\ 0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 \end{pmatrix} \begin{array}{l} \text{scale} \\ \text{orientation} \\ \text{vertical translation} \\ \text{horizontal translation} \end{array}$$

where l stands for the low frequency wavelet, h, v, d for the horizontal, vertical and diagonal orientations respectively. Figure 6 illustrates this neighbourhood \mathcal{N}_1 for a given wavelet at the center of the image.

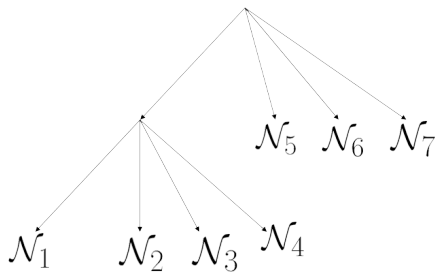


Figure 5. Quadtree corresponding to a wavelet transform of depth 2

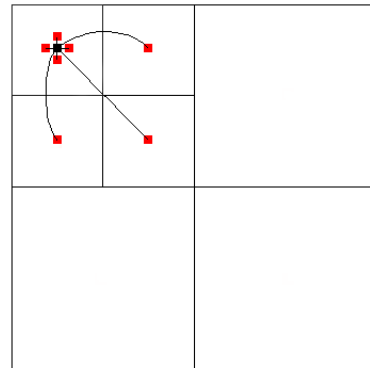


Figure 6. Illustration of the neighbourhood \mathcal{N}_1 . In black, the given wavelet at the center of the image, and in red all the preserved neighbouring wavelets.

In order to validate this notion of neighbourhood, we considered two different scenarii. The first one consists of setting for all i , $\mathcal{N}_i = \begin{pmatrix} \text{same scale} \\ \text{all orientations} \\ 0 \\ 0 \end{pmatrix}$. The second one consists of setting for all i ,

$$\mathcal{N}_i = \begin{pmatrix} \text{same scale} & \text{same scale} & \text{same scale} & \text{same scale} & \text{same scale} \\ \text{all orientations} & \text{all orientations} & \text{all orientations} & \text{all orientations} & \text{all orientations} \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 \end{pmatrix},$$

Figure 7 displays restored images using sparsity patterns generated from these two neighbourhoods. This experiment shows that sparsity patterns defined from the correlations between neighbouring wavelets are relevant for our aim: approximating and interpolating the blurring operator.

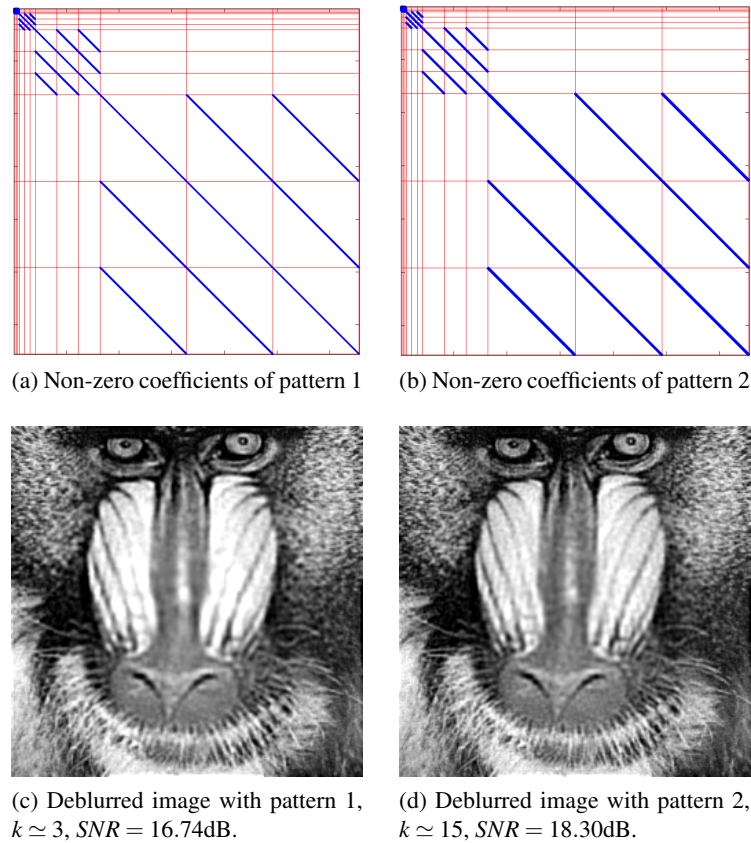


Figure 7. Deblurred images using matrices defined by sparsity patterns

Conclusion

In this paper we investigated the use of sparse matrices in the wavelet domain in order to approximate spatially varying blur operators. We observed that very large compression factors still allow to obtain near optimal reconstruction results. Compared to previously proposed approach, this technique allows to approximate the exact operator with an arbitrary precision. It is thus suited to noiseless and noisy problems. We also showed that pre-established sparsity patterns can be used to efficiently approximate the blurring integral operator. This property is central in order to tackle blind deconvolution problems. We are currently working on this topic.

References

- [1] Nagy J and O’Leary D 1998 *SIAM Journal on Scientific Computing* **19** 1063
- [2] Ben Hadj S and Blanc-Féraud L 2011 URL <http://hal.inria.fr/inria-00602650/fr/>
- [3] Zhang B, Zerubia J and Olivo-Marin J 2007 *Applied Optics* **46** 1819–1829
- [4] Hansen P, Nagy J and O’leary D 2006 *Deblurring images: matrices, spectra, and filtering* (Siam)
- [5] Escande P, Weiss P and Malgouyres F 2013 *Proceedings of ICPRAM* URL <http://hal.archives-ouvertes.fr/hal-00733194>
- [6] Beylkin G, Coifman R and Rokhlin V 1991 *Commun. Pure and Applied Math.* **44** 141–183
- [7] Coifman R and Meyer Y 1997 *Cambridge Studies in Advanced Math* **48**
- [8] Chang E C, Mallat S and Yap C 1999 *Applied and Computational Harmonic Analysis* **9** 312–335