

MAXIMUM LIKELIHOOD EXTENSION FOR NON-CIRCULANT DECONVOLUTION

Javier Portilla

Instituto de Óptica
Consejo Superior de Investigaciones Científicas (CSIC), Madrid

ABSTRACT

Directly applying circular de-convolution to real-world blurred images usually results in boundary artifacts. Classic boundary extension techniques fail to provide likely results, in terms of a circular boundary-condition observation model. Boundary reflection gives rise to non-smooth features, especially when oblique oriented features encounter the image boundaries. Tapering the boundaries of the image support, or similar strategies (like constrained diffusion), provides smoothness on the toroidal support; however this does not guarantee consistency with the spectral properties of the blur (in particular, to its zeros). Here we propose a simple, yet effective, model-derived method for extending real-world blurred images, so that they become likely in terms of a Gaussian circular boundary-condition observation model. We achieve artifact-free results, even under highly unfavorable conditions, when other methods fail.

Index Terms— non-circulant deconvolution, boundary artifacts, image restoration, maximum likelihood extension

1. INTRODUCTION

Filters used in image restoration are typically expressed as a quotient in the discrete Fourier domain, and they are usually applied in that domain. This applies both to linear and non-linear (possibly state-of-the-art) techniques, as these also typically apply iterative linear filtering, combined with some non-linear operations. By filtering through the DFT (Discrete Fourier Transform) we establish circular boundary-conditions to the observation model. In real-world degraded images this unnatural constraint, of course, does not hold. Mismatching of circular and real boundary condition typically plagues the restoration result with boundary artifacts.

In recent years, a growing interest has emerged in the boundary condition problem in image restoration. Reeves [1] considered the specific case of deconvolution being done using a regularized inverse linear operator with its classical structure (as the ones used in Richardson-Lucy, Tikhonov, Wiener filtering, etc.) to separate the (non-circulant) linear deconvolution problem as a sequence of two steps: first, to

compute a set of boundary pixels extending the observation, fulfilling some constraints (a non-translation-invariant problem, but with relatively few unknowns); and, second, to perform the filtering of the extended image, assuming circular boundary conditions. Recently Sorel [2] adapted that method to an efficient non-linear restoration technique [3]. For iterative non-linear restoration methods the modification of the original circulant deconvolution matrices also needs to be done iteratively over partial results. Other authors (e.g., [4]) have followed this same path of modifying previous restoration algorithms assuming circulant convolutions to deal with the boundary condition problem in effective, but *ad-hoc* (i.e., specific to the restoration method) ways. Here we follow the alternative approach of obtaining a single appropriated extension of the observed image, independently of the restoration algorithms (assumed DFT-based) to be applied.

Liu and Jia [5] investigated the extension, through mirror reflection, of the image to another, four times bigger in surface, providing continuity on a toroidal support (a *tiling* image, in computer graphics terminology). Then they solved for a variational problem in order to further smooth boundary transitions. Note that classical partial mirror extension does not preserve continuity on a toroidal support. In general, whenever reflection is considered, strength of artifacts will critically depend on (1) the existence of oblique features near the image boundaries (as they generate new spatial frequencies after being reflected); and (2) the lack of vertical/horizontal symmetry of the blurring kernel, which provokes that the mirrored extensions can no longer be described as filtered by the same kernel.

Another simple universal strategy to avoid boundary artifacts consists of using a spatially variant circulant smoothing operator acting only on the vicinity of the image support's boundaries, thus imposing smoothness on a toroidal support. This is implemented by the popular $\text{\textcircled{R}}$ Matlab function *edgetaper.m*. However, methods like this or the referred [5] disregard the fact that smoothness on the toroidal support, on its own, does not fully comply with the circulant observation model. It is also necessary that the extended image respects the spectral valleys of the blurring kernel. This latter observation is a central motivation of this work.

Thanks to Spanish Government funding grant TIN2012-38102-C03-01.

2. METHODS

2.1. Realistic vs. circulant observation models

A useful observation model for noisy convolved images is

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{w}, \quad (1)$$

where the \mathbf{H} matrix represents an $M^2 \times N^2$ (with $M < N$), non-circulant convolution, and \mathbf{w} is the added noise term. Whereas using a non-circulant matrix is the correct choice for real-world image convolution, most restoration algorithms assume, for convenience, a circular boundary condition observation model:

$$\mathbf{z} = \mathbf{H}_c \mathbf{x}' + \mathbf{w}', \quad (2)$$

where \mathbf{H}_c is a square $(N')^2 \times (N')^2$ block-circulant convolution matrix. We force the mean and autocovariance of the extended images to be identical to those of the original image (the latter computed/estimated with non-circular boundary conditions). A way of making compatible Eqs. (1) and (2) is to force $\mathbf{S}\mathbf{z} = \mathbf{y}$, where \mathbf{S} is a $M^2 \times (N')^2$ matrix, with $N' \geq N$, selecting the observed M^2 pixels of \mathbf{y} from its extended version. \mathbf{H}_c is a square circulant extension of \mathbf{H} . The restoration problem of non-circulant deconvolution has been thus split into three sequential steps: (1) find a *suitable* extension $\hat{\mathbf{z}}(\mathbf{y})$ of \mathbf{y} ; (2) perform a circulant deconvolution of $\hat{\mathbf{z}}(\mathbf{y})$; (3) discard the added pixels.

2.2. Problem assessment

There are $(N')^2 - M^2$ degrees of freedom for extending \mathbf{y} to \mathbf{z} . Here we follow the criterion of choosing the extension which maximizes the likelihood of an adapted, Gaussian circular boundary-condition observation model. We have assumed that both the added noise \mathbf{w} (considered white) and the original image \mathbf{x} vector terms follow an homogenous (translation-invariant) distribution. Then the Discrete Fourier Transform diagonalizes the circulant covariance matrices of their extensions. In their diagonals appear the corresponding Power Spectral Densities (PSDs):

$$\begin{aligned} P_Y(f_i) &= \mathbb{E}|Z(f_i)|^2 \\ &= |H(f_i)|^2 \mathbb{E}|X'(f_i)|^2 + \mathbb{E}|W'(f_i)|^2 \\ &= |H(f_i)|^2 P_X(f_i) + \sigma_W^2, \end{aligned} \quad (3)$$

for each of the $(N')^2$ discrete frequencies f_i , where $X'(\mathbf{f})$ represents here the Fourier transform of \mathbf{x}' (analogous with \mathbf{z} and \mathbf{w}'). Considering, for simplicity, Gaussian distributions for all the implied vectors, the observation's minus log likelihood is (ignoring constant terms) $\propto \sum_{i=1}^{N'} |Z(f_i)|^2 / P_Y(f_i)$, and our maximum likelihood (ML) extension problem is:

$$\hat{\mathbf{z}}(\mathbf{y}) = \arg \min_{\mathbf{z}} \|\mathbf{D}_{P_Y}^{-1/2} \mathbf{F}^* \mathbf{z}\|^2, \text{ s.t. } \mathbf{S}\mathbf{z} = \mathbf{y}, \quad (4)$$

where \mathbf{F}^* is a matrix performing a 2D-DFT, and \mathbf{D}_{P_Y} is a diagonal matrix whose entries correspond to $P_Y(u, v)$, both lexico-graphically ordered.

2.3. Estimating P_Y

There are two alternative paths for estimating P_Y . First one is to compute it as the DFT of the sample auto-covariance of the observation \mathbf{y} . This possibility is appealing, as it allows to attack Eq.(4) without knowing the blurring kernel h or the noise variance σ_w^2 . However, estimating sample statistics poses a trade-off between variance and bias (*noisy* vs. *blurred* estimation) whose optimal balance point may be difficult to know in advance. Second alternative consists of estimating/setting P_X first, and then applying Eq.(3). The advantage of the latter approach is that, even if our P_X estimate is rough (e.g., a generic signal model), degradation will nevertheless be reliably accounted for (provided we know h and σ_w^2 with some accuracy). In particular, it will suffice to identify the set of frequencies of the extended observation which are more heavily penalized by the cost function (low $P_Y(f_i)$). Here we have followed this second strategy.

We have chosen a Gaussian Markov Random Field model for the original image \mathbf{x} , a separable AR-1 (auto-regressive, one tap) model, whose PSD is:

$$P_X(u, v) = \frac{4\sigma_x^2 \log(\rho)^2}{(\log(\rho)^2 + 4\pi^2 u^2)(\log(\rho)^2 + 4\pi^2 v^2)}. \quad (5)$$

We have hand-optimized its parameters, obtaining $\sigma_x = 30$ and $\rho = 0.65$, for the images and experiments described in Section 3. We note that the behavior of the method was robust over change in these parameters, over a wide range.

2.4. Unconstrained reformulation

First, we split $\mathbf{z}_P = \mathbf{P}\mathbf{z} = [\mathbf{z}_e^T \mathbf{z}_i^T]^T$ into two non-overlapping vectors: \mathbf{z}_i , corresponding to the M^2 original observed pixels, and \mathbf{z}_e , corresponding to the $(N')^2 - M^2$ extended pixels. \mathbf{P} is a permutation matrix. Analogously, the Fourier transform can be re-ordered such that $\mathbf{F}_P^* = \mathbf{F}^* \mathbf{P}^{-1} = [\mathbf{F}_e^* \mathbf{F}_i^*]$, corresponding to the two component of the Fourier transform applied only to the external and the internal pixels, respectively. $\mathbf{F}^* \mathbf{z} = \mathbf{F}_P^* \mathbf{z}_P = \mathbf{F}_e^* \mathbf{z}_e + \mathbf{F}_i^* \mathbf{z}_i$, and the extension constraint implies $\mathbf{z}_i = \mathbf{y}$. Now Eq.(4) simplifies to

$$\hat{\mathbf{z}}(\mathbf{y}) = \arg \min_{\mathbf{z}_e} \|\mathbf{D}_{P_Y}^{-1/2} (\mathbf{F}_e^* \mathbf{z}_e + \mathbf{F}_i^* \mathbf{y})\|^2, \quad (6)$$

whose solution is:

$$\hat{\mathbf{z}}(\mathbf{y}) = -(\mathbf{F}_e \mathbf{D}_{P_Y}^{-1} \mathbf{F}_e^*)^{-1} \mathbf{F}_e \mathbf{D}_{P_Y}^{-1} \mathbf{F}_i^* \mathbf{y}. \quad (7)$$

In terms of practical implementation, it is useful to re-express $\mathbf{F}_e^* = \mathbf{F}^* \mathbf{E}^T$, and $\mathbf{F}_i^* = \mathbf{F}^* \mathbf{S}^T$, where \mathbf{E}^T ($((N')^2 \times ((N')^2 - M^2))$) is the extension of external pixels to a vector containing the whole image, the center being filled with zeros, \mathbf{S}^T ($((N')^2 \times M^2)$) performs analogously, but preserving the central region. Their transpose do the reverse (selection) operations. Thus, the difficulty in Eq. (7) lies only in the involved matrix inversion.

2.5. Recursive computation

Let us call $\mathbf{A} = \mathbf{F}_e \mathbf{D}_{P_Y}^{-1} \mathbf{F}_e^*$ and $\mathbf{b}(\mathbf{y}) = \mathbf{F}_e \mathbf{D}_{P_Y}^{-1} \mathbf{F}_e^* \mathbf{y}$. First note that \mathbf{A} always exists when $\sigma_w^2 > 0$. In this case existence of \mathbf{A} implies \mathbf{A}^{-1} also exists. Thus, we can use Neumann's series [6] $\mathbf{A}^{-1} = \sum_{n=0}^{\infty} (\mathbf{I} - \mathbf{A})^n$ to construct the following recursive method for computing $\hat{\mathbf{z}}(\mathbf{y}) = \mathbf{A}^{-1} \mathbf{b}(\mathbf{y})$:

$$\mathbf{z}^{(n+1)} = \mathbf{z}^{(n)} + \lambda(\mathbf{b}(\mathbf{y}) - \mathbf{A}\mathbf{z}^{(n)}),$$

with, in principle, $\lambda = 1$. However, we accelerate convergence by using $\lambda = 2/|\lambda_{\max}(\mathbf{A})|$ ($|\lambda_{\max}(\mathbf{A})|$ is the largest eigenvalue of \mathbf{A} in absolute value), without affecting the fixed-point solution. We have used a close approximation: $\lambda = 2\sigma_w^2$. As a starting point we use a mirror-extended image, edge-tapered with a Gaussian kernel (see details in next section). The observation's sample mean was subtracted, and added back after the extension.

3. EXPERIMENTS AND RESULTS

We have tested our method with three 8-bit gray-level 256^2 pixel images: two typical images (*Cameraman* and *House*), and an image from Brodatz photo texture album, (*Straw*), having a lot of energy along diagonal and oblique orientations. We have used 8 degradations, with 4 blurring kernels, each adding two levels (low and medium) of noise. We have used the following point spread functions (PSFs), all of them with normalized sum: *PSF1* is $h_{i,j} = (1 + i^2 + j^2)^{-1}$, for $i, j = -7 \dots 7$. *PSF2* is a 9×9 uniform kernel. *PSF3* simulates a uniform camera vertical movement 9×1 . *PSF4* is the oblique blurring kernel given by [0000111; 0012321; 0134310; 1232100; 1110000].

We obtained \mathbf{y} by cropping the $M \times M$ valid central region after performing circular convolution and noise addition to the original. Note that $M = N - 2L$, being $2L + 1$ the width/height of the kernel support. We have applied our extension to two very different restoration methods using circular boundary conditions: the classic Wiener filter (by using the same spectral model of Eqs. (5) and (3)), and the L_2 -relaxed L_0 method *ConDy*, optimized for 10 iterations [7]. We have compared our results, measured as increment in signal-to-noise ratio (ISNR, in decibels), to the results of: (1) Mirror extension, using $L_e = 8$ ($N' = N + 2L_e = M + 2L + 2L_e$), being L_e an extra extension of the original support (which we have tested it helps to reduce the artifacts); (2) Full mirror extension, without cropping the mirrored copies ($N' = 2M$); (3) Same boundary reflection as (1), but now followed by edge tapering using `@Matlab edgetaper.m`, with a Gaussian kernel of $\sigma = 6$ (hand-optimized value); (4) Boundary artifact-free *oracle*: obtained by restoring the whole $N \times N$ circulant convolved image, and then cropping the $M \times M$ central region.

To understand the influence of the degradation on the results we have plotted the average mean square improvement

ratio (expressed as ISNR) in Figure 1, for every simulated degradation. The *oracle* result is depicted as a white box indicating a practical upper bound performance. First, we observe that mirror extension-based methods are not robust. Limited reflection (labeled as MIRROR) only performs acceptably with *PSF1*, which is the only amongst the 4 used kernels that has no zeros in Fourier. Full reflection extension (labeled MIRROR-T, "T" for *tiling*) is far more robust, but also fails for *PSF4*. Lacking this kernel vertical/horizontal symmetry, the reflected observation is highly unlikely in terms of the observation model. In contrast, *edgetaper* rarely provokes very strong artifacts. However, its results are highly dependent on the kernel and observed image, and do not get close, on average, to the oracle performance, especially when the PSF has zeros in Fourier and the noise level is low. In such cases, the only method amongst the compared ones which behaves close to the reference oracle is our maximum likelihood extension.

Table 1 shows, for reference purposes, the complete results of our MLE method, as compared to the boundary artifact-free oracle and the *edgetaper* `@Matlab` function. There we can see, among other things, how the oblique spatial frequencies, dominant in *Straw*, affect very negatively to the *edgetaper* method, as implemented here (using mirror extension prior to the *edgetapering*). The advantage of MLE is especially prominent for the degradation experiments having low noise and PSFs having zeros in Fourier, and for those non HV-symmetric (degradation 3, 5, 7 and 8). Figure 3 shows a visual comparison of the *edgetaper* extended result vs. MLE for *Straw*, using Wiener restoration.

Finally, Figure 2 shows how the quality gap between the *edgetaper* technique and the artifact-free oracle performance is effectively saved by the MLE method. We see how 10 to 50 iterations of the algorithm suffice to obtain good results in practice. ML extension using 25 iterations in our implementation (`@Matlab` over double quad `@Intel @Xeon` architecture) takes around 0.5s for images this size.

4. CONCLUSION

Boundary artifacts in image restoration are due to the mismatch between non-circulant convolved real-life observations and the circulant observation model usually assumed by restoration methods. To overcome this mismatch, we extend the observation to make it consistent with a Gaussian circulant observation model, while keeping its auto-covariance. By maximizing likelihood on this model, we have obtained, in all studied cases, restoration results very close to those, artifact-free, obtained using circulant convolutions for simulating the observation. We have tested the method against a variety of PSFs, noise levels, original images, and even against two different restoration algorithms.

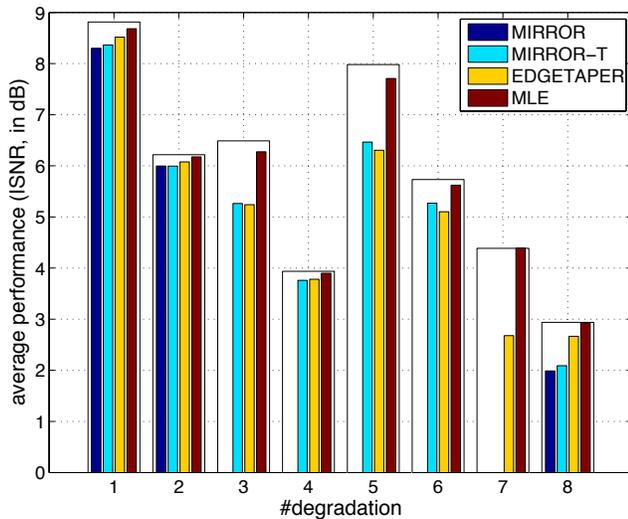


Fig. 1. Average performance, for the compared extension methods, and the 8 degradation (blur + noise) experiments.

Restor →		Wiener			ConDy10		
Method →		E-T	MLE	Ora	E-T	MLE	Ora
PSF	σ_w^2	HOUSE					
1	.25	7.41	7.63	7.82	8.55	8.86	9.15
1	2	5.01	5.12	5.16	7.51	7.80	7.90
2	.31	4.95	6.56	6.69	9.14	10.25	10.34
2	4	4.09	4.11	4.07	7.08	7.29	7.42
3	1	5.18	6.20	6.27	7.68	9.14	9.27
3	4	4.10	4.35	4.34	6.38	7.24	7.40
4	.25	2.73	3.10	3.01	4.93	5.59	5.61
4	4	1.97	2.11	2.09	4.13	4.55	4.56
PSF	σ_w^2	CAMERAMAN					
1	.25	8.67	8.70	8.71	10.56	10.58	10.62
1	2	5.56	5.55	5.56	8.06	8.11	8.12
2	.31	5.73	5.95	6.24	8.80	9.12	9.25
2	4	3.89	3.87	3.90	5.40	5.48	5.50
3	1	5.76	6.40	6.60	8.51	9.66	9.99
3	4	4.29	4.42	4.49	6.87	7.32	7.43
4	.25	3.77	3.90	3.90	6.58	6.85	6.91
4	4	2.32	2.35	2.36	3.66	3.76	3.78
PSF	σ_w^2	STRAW					
1	.25	8.97	9.26	9.39	7.63	7.72	7.82
1	2	5.85	5.97	6.01	5.33	5.41	5.46
2	.31	3.06	4.56	4.84	3.37	4.29	4.49
2	4	2.31	2.53	2.60	1.92	2.08	2.13
3	1	5.73	7.93	8.46	5.90	8.03	8.56
3	4	5.00	5.82	6.03	4.67	5.53	5.73
4	.25	2.37	4.91	4.96	-0.67	3.17	3.19
4	4	2.62	3.15	3.19	1.79	2.19	2.20

Table 1. Performance comparison, measured in Increment of Signal-to Noise Ratio, in decibels. "E-T" refers to "edgetaper", "Ora" to the oracle or upper bound reference used, and "MLE" to our ML extension method.

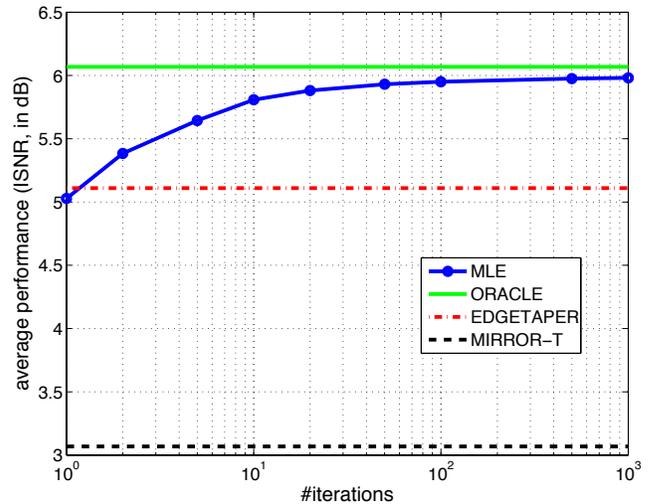


Fig. 2. Average performance, for ConDy10 restoration, as a function of the number of iterations.

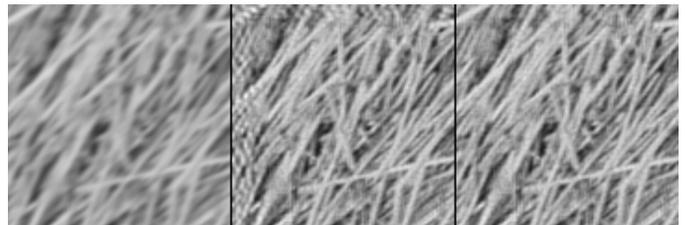


Fig. 3. Upper left corner of *Straw*. From left to right: blurred (degradation 7), E-T + Wiener, MLE + Wiener.

5. REFERENCES

- [1] S.J. Reeves, "Fast image restoration without boundary artifacts," *Image Processing, IEEE Transactions on*, vol. 14, no. 10, pp. 1448–1453, 2005.
- [2] M. Sorel, "Removing boundary artifacts for real-time iterated shrinkage deconvolution," *Image Processing, IEEE Transactions on*, vol. 21, no. 4, pp. 2329–2334, 2012.
- [3] J. Portilla, "Image restoration through ℓ_0 analysis-based sparse optimization in tight frames," in *Image Processing (ICIP), 2009 IEEE International Conference on*, 2009, pp. 3909–3912.
- [4] A. Matakos, S. Ramani, and J.A. Fessler, "Image restoration using non-circulant shift-invariant system models," in *Image Processing (ICIP), 2012 IEEE International Conference on*, 2012, pp. 3061–3064.
- [5] Renting Liu and Jiaya Jia, "Reducing boundary artifacts in image deconvolution," in *Image Processing, (ICIP), 2008 IEEE International Conference on*, 2008, pp. 505–508.
- [6] G.W. Stewart, *Matrix Algorithms: Volume 1: Basic Decompositions*, Society for Industrial and Applied Mathematics, 1998.
- [7] J. Portilla, E. Gil-Rodrigo, D. Miraut, and R. Suarez-Mesa, "ConDy: Ultra-fast high performance restoration using multi-frame L2-relaxed-L0 sparsity and constrained dynamic heuristics," in *Image Processing, (ICIP), 2011 IEEE International Conference on*, 2011, pp. 1837–1840.