Fast automatic myopic deconvolution of angiogram sequence

Louis Thibon, Ferréol Soulez, Éric Thiébaut

Fast automatic myopic deconvolution of angiogram sequences*

Louis Thibon, Ferréol Soulez and Éric Thiébaut

Centre de Recherche Astrophysique de Lyon, CNRS/UMR 5574
Université Lyon 1 – ENS Lyon, Université de Lyon, France

ABSTRACT

We present a fast unsupervised myopic deconvolution method dedicated to quasi-real time processing of video sequences such as angiograms. Our method is based on a Bayesian approach of which the tuning parameters are automatically set thanks to the marginalized likelihood of the observed image. We demonstrate the effectiveness of our approach on simulated and empirical images.

1. CONTEXT AND PREVIOUS WORK

Coronary angiography is a medical imaging technique used to visualize heart vessels (the coronaries) in order to diagnose and prevent potential heart failures. Its principle consists in injecting a radiocontrast agent by catheter into an artery and capturing an image sequence of the cardiovascular system thanks to X-ray irradiation.

In a previous work [1], we shown that the blur of angiogram sequences could be approximated by a convolution by a shift-invariant point spread function (PSF). We found that the actual PSF is constant for a given video sequence but depends on the operating conditions and on the patient. We therefore proposed to use multi-frame blind deconvolution to improve the quality of the videos both in terms of signal to noise ratio and resolution. We observed that the PSF are approximately isotropic and bell shaped with a profile similar to a Lorentzian distribution:

$$h(\boldsymbol{r}) \approx \frac{\eta}{1 + \|2\,\boldsymbol{r}/\gamma\|^2}\,, \tag{1}$$

with $\boldsymbol{r}$ the 2-D position on the detector, γ the full width at half-maximum (FWHM) of the PSF and η a normalization factor. This multi-frame blind deconvolution method is however lengthy and requires the tuning of a number of control parameters which make it impracticable for non-specialists and incompatible with the quasi real-time requirement. In this paper, we propose several improvements to achieve a fast and unsupervised method which works well in practice.

2. UNSUPERVISED MYOPIC DECONVOLUTION

2.1. Data model

For a given frame of the sequence, a raw image writes [1]:

$$\boldsymbol{y} = \mathbf{H}_\gamma \cdot \boldsymbol{x} + \boldsymbol{n} \tag{2}$$

with $\boldsymbol{y} \in \mathbb{R}^M$ the data (y_j is the value measured by j-th pixel and M is the number of pixels), $\mathbf{H}_\gamma \in \mathbb{R}^{M\times N}$ the blur operator, $\boldsymbol{x} \in \mathbb{R}^N$ the object of interest in the form of a pixelized image with N pixels and $\boldsymbol{n} \in \mathbb{R}^M$ a term which accounts for noise and approximation errors. The notation $\mathbf{H}_\gamma$ indicates that the blur depends on some parameters such as the width γ of the PSF.

2.2. Bayesian approach

Our goal is to recover the object $\boldsymbol{x}$ given the data $\boldsymbol{y}$. This is a *myopic deconvolution* problem as it requires to solve a joint problem of deblurring the data and estimating the PSF parameters. Deconvolution is well known to be an ill-conditioned inverse problem which has to be regularized to obtain a robust and stable solution [2]. The maximum a posteriori (MAP):

$$\boldsymbol{x}_{\mathsf{MAP}} = \arg\max_{\boldsymbol{x}} \Pr(\boldsymbol{x} \,|\, \boldsymbol{y}, \boldsymbol{\theta}) \tag{3}$$

provides such a solution. In words, MAP aims at maximizing the posterior probability of the unknowns $\boldsymbol{x}$ given the data $\boldsymbol{y}$ and so-called *hyper-parameters* $\boldsymbol{\theta}$ (among others the ones which characterize the PSF). Using Bayes rule:

$$\Pr(\boldsymbol{x} \,|\, \boldsymbol{y}, \boldsymbol{\theta}) = \frac{\Pr(\boldsymbol{y} \,|\, \boldsymbol{x}, \boldsymbol{\theta})\ \Pr(\boldsymbol{x} \,|\, \boldsymbol{\theta})}{\Pr(\boldsymbol{y} \,|\, \boldsymbol{\theta})}\,,$$

and realizing that the *evidence* $\Pr(\boldsymbol{y} \,|\, \boldsymbol{\theta})$ does not depend on $\boldsymbol{x}$ yields:

$$\boldsymbol{x}_{\mathsf{MAP}} = \arg\max_{\boldsymbol{x}} \Pr(\boldsymbol{y} \,|\, \boldsymbol{x}, \boldsymbol{\theta})\ \Pr(\boldsymbol{x} \,|\, \boldsymbol{\theta})\,, \tag{4}$$

where the likelihood of the data $\Pr(\boldsymbol{y} \,|\, \boldsymbol{x}, \boldsymbol{\theta})$ depends on the noise statistics and where the prior distribution $\Pr(\boldsymbol{x} \,|\, \boldsymbol{\theta})$ implements prior knowledge about the sought image $\boldsymbol{x}$. This however assumes that the hyper-parameters $\boldsymbol{\theta}$ are known. In order to have an unsupervised method, we want to be able to automatically determine $\boldsymbol{\theta}$. To that end, we propose to use maximum likelihood estimator for the hyper-parameters:

$$\boldsymbol{\theta}_{\mathsf{GML}} = \arg\max_{\boldsymbol{\theta}} \Pr(\boldsymbol{y} \,|\, \boldsymbol{\theta})\,. \tag{5}$$

Using again Bayes rule, $\Pr(\boldsymbol{y} \,|\, \boldsymbol{\theta})$ is obtained by marginalizing the joint distribution of $\boldsymbol{y}$ and $\boldsymbol{x}$ given $\boldsymbol{\theta}$:

$$\Pr(\boldsymbol{y} \,|\, \boldsymbol{\theta}) = \int \Pr(\boldsymbol{y}, \boldsymbol{x} \,|\, \boldsymbol{\theta})\,\mathrm{d}\boldsymbol{x} = \int \Pr(\boldsymbol{y} \,|\, \boldsymbol{x}, \boldsymbol{\theta})\ \Pr(\boldsymbol{x} \,|\, \boldsymbol{\theta})\,\mathrm{d}\boldsymbol{x}\,.$$

The problem now amounts to integrate the joint distribution. Owing to the number of variables (N the number of pixels in the restored image), the only realistic cases are when the integral can be carried out analytically. This implies to restrict our-self to Gaussian distributions.

*This work is part of the MiTiV project funded by the French National Research Agency (ANR DEFI 09-EMER-008-01).

2.3. Generalized maximum likelihood

We introduce the, possibly improper, Gaussian distribution of the centered variables $\boldsymbol{u}$ with symmetric positive semi-definite weighting matrix $\mathbf{W}$:

$$\mathcal{G}(\boldsymbol{u};\mathbf{W}) \stackrel{\text{def}}{=} \left|\frac{\mathbf{W}}{2\,\pi}\right|_+^{1/2} \exp\Bigl(-\frac{1}{2}\,\|\boldsymbol{u}\|_{\mathbf{W}}^2\Bigr)\,, \tag{6}$$

with $\|\boldsymbol{u}\|_{\mathbf{W}}^2 = \boldsymbol{u}^\top\cdot\mathbf{W}\cdot\boldsymbol{u}$ and $|\mathbf{W}|_+$ the product of the strictly positive eigenvalues of $\mathbf{W}$. The weighting matrix $\mathbf{W}$ is also called the *precision matrix* by some authors [3]. If the covariance matrix of $\boldsymbol{u}$ is non-singular, taking $\mathbf{W} = \mathrm{Cov}(\boldsymbol{u})^{-1}$ and noting that $|\mathbf{W}|_+$ is then the determinant of $\mathbf{W}$, the usual Gaussian distribution is retrieved. However, the definition in Eq. (6) let us cope with improper statistics.

In practice, the statistics of the error term $\boldsymbol{n}$ in Eq. (2) is well approximated by a centered Gaussian distribution, the likelihood of the data $\boldsymbol{y}$ knowing the object $\boldsymbol{x}$ and the tuning parameters $\boldsymbol{\theta}$ is then given by:

$$\Pr(\boldsymbol{y}\,|\,\boldsymbol{x},\boldsymbol{\theta}) = \mathcal{G}(\boldsymbol{y}-\mathbf{H}_\gamma\cdot\boldsymbol{x};\mathbf{B}/\sigma^2)\,, \tag{7}$$

with $\mathbf{H}_\gamma\cdot\boldsymbol{x} = \mathbb{E}(\boldsymbol{y}\,|\,\boldsymbol{x},\boldsymbol{\theta})$ and $\mathbf{B}/\sigma^2 = \mathbf{W}_{\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta}}$ the expectation and precision matrix of the data given $\boldsymbol{x}$ and $\boldsymbol{\theta}$ respectively. The scaling term σ^2 characterizes the noise level. It is, for instance, an estimate of the noise variance after some *whitening* implemented by the matrix $\mathbf{B}$ which is, at least, positive semi-definite and may be different from the identity to account for non-i.i.d. noise.

Since we restrict our-self to Gaussian distributions, the most general expression of the a priori distribution of $\boldsymbol{x}$ knowing $\boldsymbol{\theta}$ is:

$$\Pr(\boldsymbol{x}\,|\,\boldsymbol{\theta}) = \mathcal{G}(\boldsymbol{x}-\overline{\boldsymbol{x}};\mu\,\mathbf{C})\,, \tag{8}$$

with $\overline{\boldsymbol{x}} = \mathbb{E}(\boldsymbol{x}\,|\,\boldsymbol{\theta})$ and $\mu\,\mathbf{C} = \mathbf{W}_{\boldsymbol{x}|\boldsymbol{\theta}}$ the expected value and the precision matrix of $\boldsymbol{x}$ knowing $\boldsymbol{\theta}$ respectively. The matrix $\mathbf{C}$ is symmetric positive semi-definite and $\mu > 0$ is a tuning parameter. In what follows, in order to factorize the noise level characterized by σ^2, we introduce $\alpha = \sigma^2\,\mu$ as the tuning parameter for the priors.

In the considered case (affine direct model and Gaussian distributions) and providing $\mathbf{H}_\gamma^\top\cdot\mathbf{B}\cdot\mathbf{H}_\gamma + \alpha\,\mathbf{C}$ is invertible, the MAP solution is unique and has a closed form:

$$\boldsymbol{x}_{\mathsf{MAP}} = \overline{\boldsymbol{x}} + \mathbf{R}_{\boldsymbol{\theta}}\cdot(\boldsymbol{y}-\mathbf{H}_\gamma\cdot\overline{\boldsymbol{x}})\,, \tag{9}$$

with $\mathbf{R}_{\boldsymbol{\theta}}$ the reconstruction matrix given by:

$$\mathbf{R}_{\boldsymbol{\theta}} = (\mathbf{H}_\gamma^\top\cdot\mathbf{B}\cdot\mathbf{H}_\gamma + \alpha\,\mathbf{C})^{-1}\cdot\mathbf{H}_\gamma^\top\cdot\mathbf{B}\,, \tag{10}$$

which does not depend on the noise level σ^2 but solely on:

$$\boldsymbol{\theta} = \{\alpha,\gamma\} \tag{11}$$

which are thus the tuning parameters of our myopic problem: γ let us tune the PSF, while α let us set the relative weight of the priors.

Providing that $\mathbf{B}/\sigma^2 = \mathbf{W}_{\boldsymbol{y}|\boldsymbol{x},\boldsymbol{\theta}}$ does not depend on $\boldsymbol{x}$, it can be easily shown that the evidence is also a Gaussian distribution:

$$\Pr(\boldsymbol{y}\,|\,\boldsymbol{\theta}) = \mathcal{G}(\boldsymbol{y}-\mathbf{H}_\gamma\cdot\overline{\boldsymbol{x}};\mathbf{Q}_{\boldsymbol{\theta}}/\sigma^2)\,, \tag{12}$$

with $\mathbf{H}_\gamma\cdot\overline{\boldsymbol{x}} = \mathbb{E}(\boldsymbol{y}\,|\,\boldsymbol{\theta})$ the expectation of $\boldsymbol{y}$ knowing $\boldsymbol{\theta}$ and $\mathbf{Q}_{\boldsymbol{\theta}}/\sigma^2 = \mathbf{W}_{\boldsymbol{y}|\boldsymbol{\theta}}$ the corresponding precision matrix with:

$$\mathbf{Q}_{\boldsymbol{\theta}} = \mathbf{B}\cdot(\mathbf{I}-\mathbf{A}_{\boldsymbol{\theta}})\,, \tag{13}$$

where $\mathbf{I}$ is the identity and $\mathbf{A}_{\boldsymbol{\theta}} \stackrel{\text{def}}{=} \mathbf{H}_\gamma\cdot\mathbf{R}_{\boldsymbol{\theta}}$ is the so-called *influence matrix* [4].

Now the evidence $\Pr(\boldsymbol{y}\,|\,\boldsymbol{\theta})$ can be maximized with respect to σ^2 to yield the *generalized maximum likelihood* (GML) estimate of the noise variance:

$$\sigma^2_{\mathsf{GML}} = \frac{1}{M_+}\,\bigl\|\boldsymbol{y}-\mathbf{H}_\gamma\cdot\overline{\boldsymbol{x}}\bigr\|_{\mathbf{Q}_{\boldsymbol{\theta}}}^2 \tag{14}$$

with M_+ the number of strictly positive eigenvalues of $\mathbf{Q}_{\boldsymbol{\theta}}$. Replacing the noise variance by its maximum likelihood value σ^2_{GML}, maximizing the $\Pr(\boldsymbol{y}\,|\,\boldsymbol{\theta})$ with respect to the other tuning parameters amounts to minimize:

$$\mathrm{GML}(\boldsymbol{\theta}) = \frac{\bigl\|\boldsymbol{y}-\mathbf{H}_\gamma\cdot\overline{\boldsymbol{x}}\bigr\|_{\mathbf{Q}_{\boldsymbol{\theta}}}^2}{M_+\,|\mathbf{Q}_{\boldsymbol{\theta}}|_+^{1/M_+}} \tag{15}$$

which is a generalization of the criterion considered by Wahba [5] and Trouvé *et al.* [3] as it takes into account non-i.i.d. noise and the a priori expectation of the parameters $\overline{\boldsymbol{x}} = \mathbb{E}(\boldsymbol{x}\,|\,\boldsymbol{\theta})$.

2.4. Practical formulation

In our case, the noise is approximately uniform and uncorrelated but we have to cope with unmeasured pixels (see leftmost image in Fig. 1) so $\mathbf{B}$ is the diagonal matrix $\mathbf{B} = \mathrm{diag}(\boldsymbol{b})$ with $\boldsymbol{b}$ the mask of measured pixels:

$$b_j = \begin{cases} 1 & \text{if } j\text{-th pixel is valid} \\ 0 & \text{else.} \end{cases} \tag{16}$$

In order to regularize the inverse problem, we impose the smoothness of the object via its prior distribution as it is typically done in image reconstruction with $\overline{\boldsymbol{x}} = \mathbf{0}$ and:

$$\mathbf{C} = \mathbf{D}^\top\cdot\mathbf{D}\,, \tag{17}$$

where $\mathbf{D}$ is a finite difference operator. That is to say that maximizing $\Pr(\boldsymbol{x}\,|\,\boldsymbol{\theta})$ amounts to minimizing $\|\boldsymbol{x}\|_{\mathbf{C}}^2 = \|\mathbf{D}\cdot\boldsymbol{x}\|_2^2$ which is a quadratic measure of the roughness of $\boldsymbol{x}$. Using 2-D indices for the sake of clarity, our regularization typically writes:

$$\|\boldsymbol{x}\|_{\mathbf{C}}^2 = \sum_{j_1,j_2}(x_{j_1+1,j_2}-x_{j_1,j_2})^2 + \sum_{j_1,j_2}(x_{j_1,j_2+1}-x_{j_1,j_2})^2\,, \tag{18}$$

with N_1 and N_2 the restored image dimensions, thus $N = N_1\times N_2$ is the number of restored pixels.

2.5. Myopic estimation

To summarize, we propose to solve the myopic part of our problem by minimizing $\mathrm{GML}(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$. This is very like the strategy proposed by Blanc *et al.* [6]. Taking into account that $\overline{\boldsymbol{x}} = \mathbf{0}$, the GML criterion is finally given by:

$$\mathrm{GML}(\boldsymbol{\theta}) = \frac{\boldsymbol{y}^\top\cdot\mathbf{B}\cdot(\mathbf{I}-\mathbf{A}_{\boldsymbol{\theta}})\cdot\boldsymbol{y}}{M_+\,|\mathbf{B}\cdot(\mathbf{I}-\mathbf{A}_{\boldsymbol{\theta}})|_+^{1/M_+}}\,, \tag{19}$$

with M_+ the number of strictly positive eigenvalues of $\mathbf{B}\cdot(\mathbf{I}-\mathbf{A}_{\boldsymbol{\theta}})$.

For comparison, we also consider the *generalized cross validation* criterion [4] which has been extensively used as the method of choice for choosing the best hyper-parameters:

$$\mathrm{GCV}(\boldsymbol{\theta}) = \frac{\boldsymbol{y}^\top\cdot(\mathbf{I}-\mathbf{A}_{\boldsymbol{\theta}})^\top\cdot\mathbf{B}\cdot(\mathbf{I}-\mathbf{A}_{\boldsymbol{\theta}})\cdot\boldsymbol{y}}{M_+\,[\mathrm{tr}(\mathbf{I}-\mathbf{A}_{\boldsymbol{\theta}})/M_+]^2}\,. \tag{20}$$

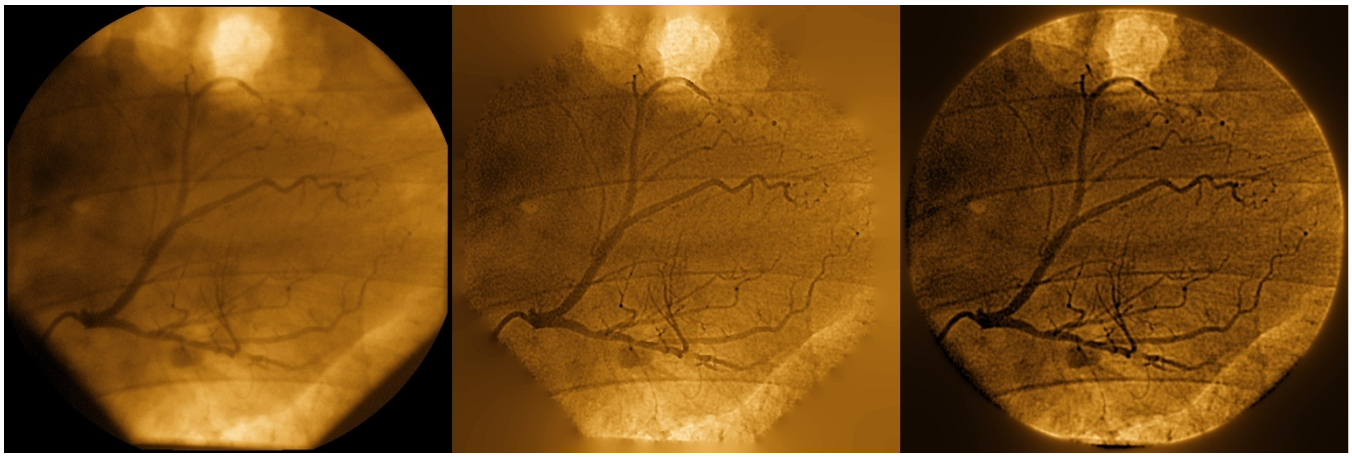

Fig. 1: Blind and myopic deconvolution of a single coronarographic image. Left: original image. Center: blind deconvolution by Soulez *et al.* [1]. Right: proposed myopic deconvolution (here with a circular mask).

2.6. Circulant approximation

The criteria (GML or GCV) which have to be minimized to determine the tuning parameters $\boldsymbol{\theta}$ are not in general applicable in large dimension because they involve to compute the determinant or the trace of a very large matrix. There exist means to approximate the quantities for large matrix [7], but our experience is that the achieved precision is insufficient for our purposes. Since $\mathbf{H}$ is a convolution, it can be approximated by a circulant block circulant matrix which is diagonalized by the discrete Fourier transform (DFT):

$$\mathbf{H}_\gamma \approx \mathbf{F}^{-1} \cdot \operatorname{diag}(\hat{\boldsymbol{h}}) \cdot \mathbf{F}\,, \tag{21}$$

with $\mathbf{F}$ the DFT operator and $\hat{\boldsymbol{h}} = \mathbf{F} \cdot \boldsymbol{h}$ the DFT of the shift invariant PSF $\boldsymbol{h}$. Note that the corresponding adjoint of the PSF operator is: $\mathbf{H}_\gamma^\top \approx \mathbf{F}^{-1} \cdot \operatorname{diag}(\hat{\boldsymbol{h}}^\star) \cdot \mathbf{F}$. The same trick can be applied to the regularization operator $\mathbf{C} = \mathbf{D}^\top \cdot \mathbf{D}$ which is also shift-invariant:

$$\mathbf{C} \approx \mathbf{F}^{-1} \cdot \operatorname{diag}(\hat{\boldsymbol{c}}) \cdot \mathbf{F}\,. \tag{22}$$

For instance, the regularization term in Eq. (18) is approximated by:

$$\|\boldsymbol{x}\|_{\mathbf{C}}^2 \approx \sum_{k1,k_2} \hat{c}_{k_1,k_2} \, |\hat{x}_{k_1,k_2}|^2\,,$$

with $\hat{\boldsymbol{x}} = \mathbf{F} \cdot \boldsymbol{x}$ the DFT of $\boldsymbol{x}$ and:

$$\hat{c}_{k_1,k_2} = 4\,\sin^2(\pi\, k_1/N_1) + 4\,\sin^2(\pi\, k_2/N_2)\,, \tag{23}$$

which is real and non-negative. Finally, by only considering a sub-image where all pixels are measured:

$$\mathbf{B} \approx \mathbf{I}\,. \tag{24}$$

Putting all together, we obtain the following approximate expressions in Fourier domain for the *generalized maximum likelihood* and *generalized cross-validation* criteria:

$$\mathrm{GML}(\boldsymbol{\theta}) \approx \frac{\sum_k \hat{q}_k \, |\hat{y}_k|^2}{\left[\prod_{k \in \mathbb{K}_+} \hat{q}_k\right]^{1/M_+}}\,, \tag{25}$$

$$\mathrm{GCV}(\boldsymbol{\theta}) \approx \frac{N \sum_k \hat{q}_k^2 \, |\hat{y}_k|^2}{\left[\sum_k \hat{q}_k\right]^2}\,, \tag{26}$$

with $\hat{\boldsymbol{h}} = \mathbf{F} \cdot \boldsymbol{y}$ the DFT of the raw image $\boldsymbol{y}$ and:

$$\hat{q}_k = \frac{\alpha\, \hat{c}_k}{|\hat{h}_k|^2 + \alpha\, \hat{c}_k}\,, \tag{27}$$

and where $\mathbb{K}_+$ is the set of frequels for which $\hat{q}_k \neq 0$ and M_+ is the size of this set.

2.7. Algorithm summary

To process a sequence of images, we first estimate the hyper-parameters on a sub-region of a selected image of the sequence by minimizing $\mathrm{GML}(\boldsymbol{\theta})$ under circulant approximation. To carry out this optimization, we use Nelder-Mead *simplex* algorithm [8] as described in [9] with simple modifications to implement bound constraints on the parameters $\boldsymbol{\theta}$. Given the GML estimate of the hyper-parameters, we perform a regularized deconvolution for each image of the sequence. This amount to solve the linear equations:

$$(\mathbf{H}_\gamma^\top \cdot \mathbf{B} \cdot \mathbf{H}_\gamma + \alpha\, \mathbf{C}) \cdot \boldsymbol{x} = \mathbf{H}_\gamma^\top \cdot \mathbf{B} \cdot \boldsymbol{y}\,. \tag{28}$$

At least due to the non-stationary operator $\mathbf{B}$, these equations cannot be easily diagonalized by FFT, so we use an iterative method such as the linear conjugate gradients [10].

3. RESULTS

We tested our method on simulated images and on a set of empirical angiogram sequences provided by the cardiac center of Hôpital de la Croix Rousse (Lyon, France). Figures 1 and 2 show typical results for empirical and simulated images respectively.

3.1. Validation of our approximations

In practice, the tuning parameters obtained from the minimization with respect to α and γ of the two versions of the GML criterion, the exact one in Eq. (19) and the approximated one in Eq. (25), are very similar. Following Trouvé *et al.* [3], we used the *generalized singular value decomposition* (GSVD) of the operators $\mathbf{H}$ and $\mathbf{D}$

Fig. 2: Myopic deconvolution of an image blurred by a Gaussian with FWHM $\gamma = 2$ pixels and with $\mathrm{SNR} = 8\,\mathrm{dB}$. *A*: result with finite difference regularization, *cf.* Eq. (23), and GML parameter estimation. *B*: result following method in [11]. *C*: result with same regularization as in *B*, *cf.* Eq. (29), but with GML parameter estimation.

to speed-up the computations of the exact criterion. The DFT approximation is however much faster[1], so we favor the use of our DFT-based method.

3.2. Comparison with other approaches

When the problem was solely to estimate the regularization level α, Wahba [5] found that GML and GCV give equally good results. However, in our myopic context where both α and the width γ of the PSF are unknown, we observed that GML provides much better estimates of the tuning parameters than GCV. This confirms the conclusions of [3] who also favored GML rather than GCV to estimate the depth from defocus of real images.

Compared to the original method by Soulez *et al.*[1], our method gives similar results in terms of quality (see Fig. 1) in spite of the simplifications made (myopic instead of blind deconvolution and quadratic regularization instead of total variation). The myopic approach is however much faster: on an Intel Core i7 CPU at 2.94 GHz, it takes about 0.4 second per 512×512 image instead of 80 seconds for the blind method. Including the myopic estimation step, a typical angiogram sequence of 80 images is processed in only 30 seconds. Our myopic algorithm is therefore a quasi real-time unsupervised method.

Finally, we also compared our strategy to the deterministic approach of [11]. They considered a quadratic regularization such that:

$$\hat{c}_{k_1,k_2} = \frac{1}{1 + \left(\sqrt{k_1^2 + k_2^2}/k_0\right)^p}\,, \tag{29}$$

where k_0 is the cut-off frequency (in frequels) of the object and p is the exponent of the decreasing rate. Figure 2 shows that their regularization which requires to fit 2 more parameters (k_0 and p) and have therefore more flexibility to adapt to the actual object structure yields an improved restoration compared to our simple regularization, given in Eq. (18). However, this is only true if parameter estimation is done via GML, not by fitting the power spectrum of the data as suggested by [11].

4. CONCLUSION

We have presented a fast unsupervised algorithm for myopic deconvolution of images. Our algorithm was tested on simulations and empirical sequences of medical images (angiograms). The first stage of the method consists in estimating the tuning parameters by means of the *generalized maximum likelihood* (GML) while the second stage is a regularized deconvolution. To estimate quickly the GML criterion, we rely on circulant approximations of the operators which are thus diagonalizable by discrete Fourier transform (DFT). We have validated our assumptions on simulations and by comparing the approximated criterion to the exact one.

Our algorithm is currently under testing at the Hôpital de la Croix Rousse (Lyon, France). Processing a typical coronarographic sequence of 80 images with 512×512 pixels takes only 30 seconds. This is about 4 times the observing duration. Thus our algorithm can be used right after the acquisition to quickly enhance the quality of the image and help the physician to make an improved diagnostic. A graphical interface have been developed and the software will be made freely available at `http://mitiv.univ-lyon1.fr`.

[1] Solving the exact myopic problem by GSVD can take up to 2 hours of computation on a 16×16 region of interest; compared to a few milliseconds when using the DFT, even on a larger, say 64×64, region.

5. REFERENCES

[1] F. Soulez, É. Thiébaut, Y. Tourneur, A. Gressard, and R. Dauphin, "Blind deconvolution of video sequences," in *15th International Conference on Image Processing*, San Diego, USA, 2008, ICIP. 1, 3, 4

[2] D. M. Titterington, "General structure of regularization procedures in image reconstruction," vol. 144, pp. 381–387, 1985. 1

[3] P. Trouvé, F. Champagnat, G. le Besnerais, and J. Idier, "Single image local blur identification," in *18th IEEE International Conference on Image Processing*. Sep 2011, pp. 621–624, Institute of Electrical and Electronics Engineers. 2, 3, 4

[4] P. Craven and G. Wahba, "Smoothing noisy data with spline functions," *Numerische Mathematik*, vol. 31, pp. 377–403, 1979. 2

[5] G. Wahba, "A comparison of GCV and GML for choosing the smoothing parameter in the generalized spline smoothing problem," *Annals of Statistics*, vol. 13, no. 4, pp. 1378–1402, 1985. 2, 4

[6] A. Blanc, L.M. Mugnier, and J. Idier, "Marginal estimation of aberrations and image restoration by use of phase diversity," vol. 20, pp. 1035–1045, 2003. 2

[7] D. Girard, "A fast 'monte-carlo cross-validation' procedure for large least squares problems with noisy data," *Numr. Math.*, vol. 56, pp. 1–23, 1989. 3

[8] J. Nelder and R. Mead, "A simplex method for function minimization," *Comp. J.*, vol. 7, pp. 308–313, 1965. 3

[9] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the Nelder–Mead simplex method in low dimensions," vol. 9, no. 1, pp. 112–147, 1998. 3

[10] M. Hestenes and E. Stiefel, "Methods of conjugate gradients for solving linear systems," *Journal of Research of the National Bureau of Standards*, vol. 49, no. 6, pp. 409–436, 1952. 3

[11] D. Gratadour, D. Rouan, L.M. Mugnier, T. Fusco, Y. Clénet, E. Gendron, and F. Lacombe, "Near-infrared adaptive optics dissection of the core of ngc 1068 with naos-conica," vol. 446, pp. 813–825, 2006. 4