

UCLA
COMPUTATIONAL AND APPLIED MATHEMATICS

Total Variation Based Oversampling of Noisy Images

Francois Malgouyres

December 2000

CAM Report 00-40

Department of Mathematics
University of California, Los Angeles
Los Angeles, CA. 90095-1555

<http://www.math.ucla.edu/applied/cam/index.html>

Total variation based oversampling of noisy images

François Malgouyres
 Dept of Mathematics
 UCLA, 6363 Math science building
 Los Angeles, CA 90095-1555, USA

Abstract

We propose and argue, in this paper, a variational model which permits to simultaneously deblur and oversample an image. Indeed, after some recalls on an existing variational model for image oversampling, we show how to modify it in order to properly achieve our two goals. We argue the modification both under the theoretical point of view of the analysis of the preservation of some structural elements and under the practical point of view of experimental results. We also describe the algorithm we use to compute a solution of this model.

1. Introduction

This paper deals with a variational methods whose aim is to both deblur and oversample an image. More precisely, noting \mathbb{T} the torus, we expect, for $N \in \mathbb{N}$, to recover an image $v \in L^2((N\mathbb{T})^2)$, from a data $u \in \mathbb{R}^{N^2}$, given by

$$u_{m,n} = (s * v)(m, n) + b_{m,n}$$

where $s \in L^2((N\mathbb{T})^2)$, $(m, n) \in \{0, \dots, N-1\}^2$ and $b \in \mathbb{R}^{N^2}$ is a Gaussian noise. Note that here the convolution is between two functions of $(N\mathbb{T})^2$. For commodity, in the following, we will denote by \mathbb{T}_N the periodization of $\{0, \dots, N-1\}$.

A very useful framework for this kind of problem is the Fourier domain. Therefore, we define the Fourier transform of a function $v \in L^2((N\mathbb{T})^2)$ by

$$\hat{v}_{\frac{\xi}{N}, \frac{\eta}{N}} = \int_{(N\mathbb{T})^2} v(x, y) e^{-2i\pi \frac{(\xi x + \eta y)}{N}} dx dy,$$

for $(\xi, \eta) \in \mathbb{Z}^2$ and the discrete Fourier transform of $u \in \mathbb{R}^{N^2}$ by

$$\hat{u}_{\xi, \eta} = \sum_{m, n=0}^{N-1} u_{m, n} e^{-2i\pi \frac{\xi m + \eta n}{N}},$$

for $(\xi, \eta) \in \{-\frac{N}{2} + 1, \dots, \frac{N}{2}\}^2$.

Using the Poisson formula (see [10], pp 29), we can express the discrete Fourier transform of u in terms of the discrete Fourier transform of b and the Fourier transform s, v . This gives

$$\hat{u}_{\xi, \eta} = \sum_{k, l \in \mathbb{Z}} \hat{s}_{\frac{\xi}{N} + k, \frac{\eta}{N} + l} \hat{v}_{\frac{\xi}{N} + k, \frac{\eta}{N} + l} + \hat{b}_{\xi, \eta},$$

for any $(\xi, \eta) \in \{-\frac{N}{2} + 1, \dots, \frac{N}{2}\}$. Therefore, we can remark that if the function v satisfies, for any $(\xi, \eta) \in \{-\frac{N}{2} + 1, \dots, \frac{N}{2}\}$, $v_{\frac{\xi}{N} + k, \frac{\eta}{N} + l} = 0$ for $k \neq 0$ or $l \neq 0$, this would be a deblurring problem. On the other hand, if we do not take into account the blurring and the noise but try to extrapolate the high frequencies, this is an interpolation problem. However, we do believe that these two issues cannot be separated and should be treated simultaneously. This is what we will propose in what follows.

There are only few paper which deals with the possibility to simultaneously deblur and oversample an image while there is an extensive literature on both deblurring and image oversampling. Therefore, concerning image deblurring, the reader can refer to [3] for most of the linear methods, to [5, 12] for variational ones (respectively based on a regularization with the entropy and the total variation) and to [11, 8] for wavelet packet based methods. Concerning oversampling, most of the linear methods tend to compute or approximate the sinc-interpolation (see [14, 15]). Non-linear methods often try to adapt the filter to the particular behavior of the image (edge, smooth region,...) (see [2, 13]) or use a regularization approach [7, 8].

The paper is organized as follows, in Section 2, we will make some recalls on a variational oversampling method introduced in [9]. Then, in Section 3, we show how to adapt this model in order to take into account the noise b . We show also that, with regard to the analysis of the preservation of a family of structural elements, this model has to be modified. We then explain the numerical scheme which is used to compute a solution of this model. Finally, in Section 4, we present some numerical experiments which confirm the importance of the modifications we have proposed in Section 3.

2 Variational oversampling of noise free images

All the results announced in this section are rigorously stated and proved in [9]. In this paper, we studied the possibility to oversample images by mean of a Maximum A Posteriori model. More precisely, we studied a variational oversampling method based on the minimization of the total variation. This method consists in finding an image $w \in L^2((N\mathbb{T})^2)$ which

$$\text{minimizes } \int_{(N\mathbb{T})^2} |\nabla w|, \text{ among } w \in \mathcal{W}_{s,u}, \quad (1)$$

where, for any given data $u \in \mathbb{R}^{N^2}$ and any convolution kernel $s \in L^2((N\mathbb{T})^2)$, we define

$$\mathcal{W}_{s,u} = \left\{ w \in L^2((N\mathbb{T})^2), \forall (m,n) \in \mathbb{T}_N^2, \right. \\ \left. s * w(m,n) = u_{m,n} \right\}.$$

Remark that, for simplicity, we note the total variation $\int_{\mathbb{T}^2} |\nabla w|$ instead of $|Dw|(\mathbb{T}^2)$.

We know that (1) has a solution as long as s is such that $\mathcal{W}_{s,u}$ is not empty and all the elements of $\mathcal{W}_{s,u}$ have the same mean. We cannot guaranty the uniqueness of this solution. However, we are sure that two different solutions have locally the same level lines at locations where these latter are properly defined and the solutions are C^1 . We also know a discretization of (1) which permits to properly approximate its solution. All these mathematical properties guaranty that this problem is well posed.

However, a drawback of this model is that if s is too much localized in space domain some points are not enough constrained by the data fidelity term. For instance if $s = \delta$ (the Dirac delta function), only the points of \mathbb{T}_N^2 suffer from the constraint of (1) and, since \mathbb{T}_N^2 is of measure zero in $(N\mathbb{T})^2$, the solutions of (1) are the constant functions¹. Therefore, arises the question of knowing whether s enough spread the constraint over the whole domain or not.

We can give an answer to that question by investigating the preservation of structural elements that are the ‘‘cylindrical functions’’. These functions are basically 1D functions and was first introduced to model whether an oversampling methods properly restore edges or not. They are rigorously defined by

Definition 1 Let $u \in \mathbb{R}^{N^2}$ and $(\alpha, \beta) \in \mathbb{Z}^2 \setminus \{(0,0)\}$. u is cylindrical along the direction (α, β) if and only if its

¹Note that here we cheat a little since when $s = \delta$ all the elements of $\mathcal{W}_{s,u}$ do not have the same mean and we are therefore not sure of the existence of a solution. Note also that this is coherent with the fact that constant functions (which are the ‘‘solutions’’ of (1) in this case) do, a priori, not belong to $\mathcal{W}_{s,u}$.




Figure 1: The hard line represents the spectral support of a function cylindrical in the direction $(2, 1)$. The thin line represents the support of the spectrum of the oversampling of this image by mean of (1).

Discrete Fourier Transform is supported by

$$\left\{ (k, l) \in \left\{ -\frac{N}{2} + 1, \dots, \frac{N}{2} \right\}^2, \beta k - \alpha l = 0 \right\}.$$

Given this definition, we can state

Proposition 1 Let N be an integer, $(\alpha, \beta) \in \mathbb{R}^2 \setminus \{(0,0)\}$, $u \in \mathbb{R}^{N^2}$ cylindrical along the direction (α, β) . For any kernel $s \in L^2((N\mathbb{T})^2)$, such that

$$\hat{s}_{\frac{\xi}{N}, \frac{\eta}{N}} \begin{cases} \neq 0 & , \text{ for } (\xi, \eta) \in \left\{ -\frac{N}{2} + 1, \dots, \frac{N}{2} \right\}^2 \\ = 0 & , \text{ otherwise,} \end{cases}$$

(1) admits a solution cylindrical along the same direction (α, β) .

Remark that this proposition can be extended to other kinds of convex regularity criterion such that the reformulation of (1) with this new regularity criterion has a solution. However, the advantage of the total variation is that it allows to reconstruct some high frequencies in a more complex manner than just filling them with 0. The strong and thin lines on Figure 1 represent the spectrum of a variational oversampling such as the one defined by (1). Once again, this has already been argued in [9] where one can find lots of experiments on this fact. However, we display on Figure 3 and 4 two experiments which show the practical importance of this proposition.

We are conscious of the fact that there could be more precise conditions on s in order to spread the constraint over $N\mathbb{T}$ (we indeed do not have any theoretical argument to assert that the condition given in Proposition 1 is necessary or optimal).

3 Variation oversampling of noisy images

3.1 The model

The model of the preceding section does not take into account the corruption of the image by noise. This has led us to introduce a model where the constraint is weaker than in (1). Therefore, the following model is very close to Rudin-Osher-Fatemi functional (see [12]) for the deblurring issue except that, this time, we do take into account the sampling and the aliasing as being part of the degradation process. This modification has already been evoked but not fully explored in [8].

More precisely an adaptation of Rudin-Osher-Fatemi method which would take into account the sampling process could be the minimization, among² $w \in BV((N\mathbb{T})^2)$,

$$\int_{(N\mathbb{T})^2} |\nabla w| + \lambda \sum_{m,n=0}^{N-1} |s * w(m, n) - u_{m,n}|^2, \quad (2)$$

for a parameter $\lambda > 0$.

However, similarly to previously, an issue of this model is that data fidelity term may not be spread enough on the whole torus (it may be concentrated in the vicinity of integer points). We can, however, state a result similar to Proposition 1 for this model which suggests to modify (2) in such a way that we avoid this artifact.

Proposition 2 *Let N be an integer, $(\alpha, \beta) \in \mathbb{R}^2 \setminus \{(0, 0)\}$, $u \in \mathbb{R}^{N^2}$ cylindrical along the direction (α, β) . For any kernel $s \in L^2((N\mathbb{T})^2)$, such that*

$$\hat{s}_{\frac{\xi}{N}, \frac{\eta}{N}} \begin{cases} \neq 0 & , \text{for } (\xi, \eta) \in \{-\frac{N}{2} + 1, \dots, \frac{N}{2}\}^2 \\ = 0 & , \text{otherwise,} \end{cases}$$

the minimization of (2) admits a solution cylindrical along the same direction (α, β) .

The proof of this proposition is very close to the one of Proposition 1 (see [8]).

Therefore, we have modified (2), according to this proposition and have minimized

$$\int_{(N\mathbb{T})^2} |\nabla w| + \lambda \sum_{m,n=0}^{N-1} |\tilde{s} * w(m, n) - u_{m,n}|^2, \quad (3)$$

²One can refer to [6] for a definition of $BV((N\mathbb{T})^2)$. However, here, it has to be understood as the space of the functions $w \in L^2((N\mathbb{T})^2)$, such that $\int_{(N\mathbb{T})^2} |\nabla w| < \infty$.

where λ is a parameter and \tilde{s} is defined by

$$\hat{\tilde{s}}_{\frac{\xi}{N}, \frac{\eta}{N}} = \begin{cases} \hat{s}_{\frac{\xi}{N}, \frac{\eta}{N}} & , \text{if } (\xi, \eta) \in \{-\frac{N}{2} + 1, \dots, \frac{N}{2}\}^2 \\ 0 & , \text{otherwise.} \end{cases}$$

In order to understand the consequence of this change, we have expressed, using Poisson formula, the data fidelity term of (2) in frequency domain and found that

$$\begin{aligned} & \sum_{m,n=0}^{N-1} |s * w(m, n) - u_{m,n}|^2 \\ &= \sum_{\xi, \eta=0}^{N-1} \left| \sum_{k,l \in \mathbb{Z}} \hat{s}_{\frac{\xi}{N} + k, \frac{\eta}{N} + l} \hat{w}_{\frac{\xi}{N} + k, \frac{\eta}{N} + l} - \hat{u}_{\xi, \eta} \right|^2. \quad (4) \end{aligned}$$

Therefore, any change in the repartition of $\hat{u}_{\xi, \eta}$ over $(w_{\frac{\xi}{N} + k, \frac{\eta}{N} + l})_{k,l \in \mathbb{Z}}$, such that $\sum_{k,l \in \mathbb{Z}} \hat{s}_{\frac{\xi}{N} + k, \frac{\eta}{N} + l} \hat{w}_{\frac{\xi}{N} + k, \frac{\eta}{N} + l}$ remains unchanged, yields the same values for this data fidelity term. Therefore, the minimization of (2) may spread $\hat{u}_{\xi, \eta}$ over these coefficient. This yields, in space domain, to a result which look like a sum of functions which are almost Dirac delta functions.

On the other hand, a formula similar to (4), for the model (3), shows that the data fidelity term deals only with low frequencies (the one in $\{-\frac{1}{2} + \frac{1}{N}, \dots, \frac{1}{2}\}^2$). Therefore, we are sure to preserve the main aspects of the image. Moreover, the heuristic of this model is to consider the aliasing as a noise.

3.2 Numerical implementation

In Section 4, we present some images which are solutions of (2) and (3). These solutions are computed using the same and more general algorithm: one which minimizes (2) for an arbitrary kernel s . Let us describe this algorithm.

The first issue in order to minimize (2) is to discretize it. Therefore, we have only considered oversampling of level $K \geq 2$. This means that in practice the result is simply an array of size $KN \times KN$. Moreover, we have discretize ∇w by a simple finite difference scheme and defined the partial derivatives

$$\Delta^x w_{m,n} = w_{m+1,n} - w_{m,n} \text{ and } \Delta^y w_{m,n} = w_{m,n+1} - w_{m,n}.$$

Moreover, in order to have a proper descent direction, we have replaced the total variation by

$$\sum_{i,j=0}^{KN-1} \sqrt{\beta^2 + (\Delta^x w_{m,n})^2 + (\Delta^y w_{m,n})^2},$$

for $\beta \in \mathbb{R}$. We call E_β the sum of this term and of the data fidelity term. Note that, in practice, we let β decrease to 0

during the iteration process. These ideas are now classical and are already discussed in [1, 4].

As we said previously, the main difference with the usual minimization of Rudin-Osher-Fatemi functional is that the data fidelity term now takes into account the sampling process. The computation of the data fidelity term and of its gradient are in this case simpler in Fourier domain. Therefore, we express it by an adaptation of (4).

More precisely, (4) becomes now

$$\begin{aligned} & \sum_{m,n=0}^{N-1} |s * w(m, n) - u_{m,n}|^2 \\ &= \sum_{\xi,\eta=0}^{N-1} \left| \sum_{k,l=0}^{K-1} \hat{s}_{\xi+kN,\eta+lN} \hat{w}_{\xi+kN,\eta+lN} - \hat{u}_{\xi,\eta} \right|^2, \end{aligned}$$

where the hats denote either the discrete Fourier transform of a signal of size $N \times N$ or $KN \times KN$ (and are therefore periodic of period N or KN).

Therefore, a simple computation permits to find that the Fourier transform of the gradient of this term is, at the frequency $(\xi + kN, \eta + lN)$, for $(\xi, \eta) \in \{0, \dots, N-1\}$ and $(k, l) \in \{0, \dots, K-1\}$,

$$2 \bar{\hat{s}}_{\xi+kN,\eta+lN} \left(\sum_{k',l'=0}^{K-1} \hat{s}_{\xi+k'N,\eta+l'N} \hat{w}_{\xi+k'N,\eta+l'N} - \hat{u}_{\xi,\eta} \right)$$

where the bar denotes the complex conjugate.

Finally, the discretization of the functional (2) is minimized with a gradient descent algorithm with an optimal step. More precisely, we start from the zero-padding (or sinc interpolation, see [15]) oversampling of the blurred image and call it u^0 . To get u^{j+1} from u^j , for $j \in \mathbb{N}$, the gradient of the functional at u^j , $\nabla E_\beta(u^j)$, is calculated, as explained above, at each step of the algorithm. Then the optimal amplitude of the variation of the image in the direction $-\nabla E_\beta(u^j)$ is estimated by the resolution of

$$\min_{s>0} E_\beta(u^j - s \nabla E_\beta(u^j))$$

using a dichotomy method. Once the optimal amplitude s_0 is calculated, we let

$$u^{j+1} = u^j - s_0 \nabla E_\beta(u^j).$$

We then iterate this process.

Note that in order to increase the speed of this algorithm, we can start from a deblurred version of u^0 , instead of u^0 . Moreover, it is better to start with a large β and to let it decrease to 0.



Figure 2: Reference image. That is the image we use for all the experiments of Section 4.

4 Experiments

All the images presented here come from manipulations (degradations and reconstructions) of the image displayed in Figure 2.

4.1 The noise free case

We present in this section some experiments which show to evidence the relevance of our interpretation of Proposition 2 when there is no noise. Remark that in such a case, the choice of the parameter λ in (2) and (3) is arbitrary as long as λ is large enough and the sampled image does not contain too much aliasing.

On Figure 3, we display some extracted part of

- Up-Left: Reference image.
- Up-Right: the downsampled image (with $s = \delta$) of Reference image (without noise).
- Down-Left: an oversampling of the downsampled image by mean of (2), with $s = \delta$ and $\lambda = 10$. Since in (2) we took $s = \delta$, we clearly see on this image the points of the grid which are constrained. Note that this drawback is still present for all the values of λ we have tested.
- Down-Right: an oversampling of the downsampled image by mean of (3), with $\lambda = 10$. Note that the points which were visible on the previous image are no longer present on this image.




Figure 3: Up-Left: the initial image. Up-Right: the down-sampling of the initial image with $s = \delta$. Down-Left: oversampling by minimizing (2). Down-Right: oversampling by minimizing (3).

On Figure 4, we display

- Up-Left: Reference image.
- Up-Right: the downsampled image (with $s = 1_{|[-\frac{1}{2}, \frac{1}{2}]^2}$) of Reference image (without noise).
- Down-Left: an oversampling of the downsampled image by mean of (2), with $s = 1_{|[-\frac{1}{2}, \frac{1}{2}]^2}$ and $\lambda = 10$. The texture is distorted. Note that the only way to avoid this distortion is to remove the texture. This depends when we take $\lambda \leq 0.1$.
- Down-Right: an oversampling of the downsampled image by mean of (3), with $\lambda = 10$. This time the texture is preserved. This illustrates the interest of Proposition 2 and of the modification introduced in (3).

Remark that this latter experiment shows that even when the sequence $(s * w(m, n))_{(m, n) \in \mathbb{T}_N^2}$ depends on all the values³ of $w(x, y)$, the modification suggested by Proposition 2 still permits to improve the method.

4.2 The noisy case

Let us now investigate the possibility to remove a potential noise while oversampling the image.

³More precisely, there does not exist any open set Ω such that, for any arbitrary modification of w on Ω , $(s * w(m, n))_{(m, n) \in \mathbb{T}_N^2}$ remains unchanged




Figure 4: Up-Left: the initial image. Up-Right: the down-sampling of the initial image with $s = 1_{|[-\frac{1}{2}, \frac{1}{2}]^2}$. Down-Left: oversampling by minimizing (2). Down-Right: oversampling by minimizing (3).




Figure 5: Up-Left: the initial image. Up-Right: the down-sampling of the initial image with $s = 1_{|[-\frac{1}{2}, \frac{1}{2}]^2}$ plus a Gaussian noise of standard deviation 3. Down-Left: oversampling by minimizing (2). Down-Right: oversampling by minimizing (3).

In order to test our algorithm, we have made the same experiment as for the creation of Figure 4 (the one with $s = 1_{|[-\frac{1}{2}, \frac{1}{2}]^2}$) except that we added a Gaussian noise of standard deviation 3. We have displayed the result of this experiment on Figure 5. For both reconstructed images the parameter λ has been fixed in such a way that the amount of remaining noise is reasonable in an homogeneous region (see Figure 6 where a larger part of the reconstruction by mean of (3) is displayed). We have taken $\lambda = 0.5$ for the oversampling by mean of (2) and $\lambda = 0.3$ for the one which uses (3). Note that the fact that for a comparable amount of noise we need a smaller λ with model (3) is not surprising since in this model the data fidelity term constrained more the image than in (2). The main comment on these images is that, despite the noise, our interpretation of Proposition 2 still makes sense.

We have also compared our results to the one obtained by a simple combination of linear algorithms. Once again, we compare them for the downsampling of the image displayed on Figure 2 with $s = 1_{|[-\frac{1}{2}, \frac{1}{2}]^2}$ and a Gaussian noise of standard deviation 3. Therefore, we display on Figure 6, on left, an oversampling obtained by composing a wiener filter (see [3]) applied to the sampled image, to deblur and denoise it, and a sinc-interpolation (see [15]), in order to oversample it (Up: the result, Down: its spectrum). We have tried to fix the parameter of the wiener filter (the assumed standard deviation of the noise) in such a way that we have the same amount of noise on this image as on the oversampling by mean of (3). However, it was not possible to find such a value for the parameter without removing all the information contained in the image. Therefore, the left images displayed on Figure 6 correspond to a value of $\sigma = 20$ and still contain a significant amount of noise.

The images on the right hand side of Figure 6 correspond to the result and its spectrum when minimizing (3) for a parameter $\lambda = 0.3$. We clearly see here that (3) permits to obtain a result which contains less noise and is sharper than with the previous method. This is also visible on the spectrum of these images. In the first case, we have just fill the high frequencies with zero and with (3), we have rebuilt a realistic spectrum out of the initial spectral domain.

Acknowledgments

I would like to thank F. Guichard, J.M. Morel and S. Durand for all the fruitful discussion we had on this subject. I would also like to thank UCLA which has financed this work.

References

- [1] R. Acar and C. Vogel. Analysis of bounded variation methods for ill-posed problems. *Inverse problems*, 10:1217–1229, 1994.
- [2] J. Allebach and P. W. Wong. Edge-directed interpolation. In *Proceedings of the International Conference on Image Processing*, volume 3, pages 707–710, 1996.
- [3] H.C. Andrews and B.R. Hunt. *Digital signal processing*. Technical Englewood Cliffs, NJ: Prentice-Hall, 1977.
- [4] T. F. Chan and P. Mulet. On the convergence of the lagged diffusivity fixed method in total variation image restoration. *SIAM Journal of Numerical Analysis*, 36(2):354–367, 1999.
- [5] G. Demoment. Image reconstruction and restoration: Overview of common estimation structures and problems. *IEEE Transactions on acoustics, speech and signal processing*, pages 2024–2036, 1989.
- [6] L.C. Evans and R. F. Gariepy. *Measure Theory and fine properties of functions*. Studies in Advanced Mathematics. CRC Press, Boca Raton, 1992.
- [7] N. B. Karayiannis and A. N. Venetsnopoulos. Image interpolation based on variational principles. *Signal Processing*, 26(1):259–288, January 1992.
- [8] F. Malgouyres. *Increase in the resolution of digital images: Variational theory and applications*. PhD thesis, ENS Cachan, 2000. available at <http://www.math.ucla.edu/~malgouy>.
- [9] F. Malgouyres and F. Guichard. Edge direction preserving image zooming: a mathematical and numerical analysis. *SIAM, J. Num. Anal.*, 2000. To appear, a preliminary version is available at: <http://www.math.ucla.edu/~malgouy>.
- [10] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, Boston, 1998.
- [11] B. Rougé. Fixed chosen noise restauration (fcnr). In *IEEE 95 Philadelphia*, 1995.
- [12] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
- [13] S. Thurnhofer and S. K. Mitra. Edge-enhanced image zooming. *Optical Engineering*, 35(7):1862–1870, July 1996.
- [14] M. Unser, A. Aldroubi, and M. Eden. Enlargement or reduction of digital image with minimum loss of information. *IEEE Transactions on Image Processing*, 4(3):247–258, March 1995.
- [15] L. Yaroslavsky. Efficient algorithm for discrete sinc interpolation. *Applied Optics*, 36(2):460–463, 1997.




Figure 6: Left: oversampling with a linear filter (wiener filter + sinc-interpolation) (Up: the image, Down: its spectrum). Right: oversampling by mean of (3) (Up: the image, Down: its spectrum).