



HAL
open science

Non-local Regularization of Inverse Problems

Gabriel Peyré, Sébastien Bogleux, Laurent D. Cohen

► **To cite this version:**

Gabriel Peyré, Sébastien Bogleux, Laurent D. Cohen. Non-local Regularization of Inverse Problems. European Conference on Computer Vision (ECCV'08), Oct 2008, Marseilles, France. pp.57-68. hal-00365030

HAL Id: hal-00365030

<https://hal.science/hal-00365030>

Submitted on 2 Mar 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Non-local Regularization of Inverse Problems

Sébastien Bogleux, Gabriel Peyré and Laurent Cohen

Ceremade, Université Paris-Dauphine, 75775 Paris Cedex 16, France
{peyre,bogleux,cohen}@ceremade.dauphine.fr

Abstract. This article proposes a new framework to regularize linear inverse problems using the total variation on non-local graphs. This non-local graph allows to adapt the penalization to the geometry of the underlying function to recover. A fast algorithm computes iteratively both the solution of the regularization process and the non-local graph adapted to this solution. We show numerical applications of this method to the resolution of image processing inverse problems such as inpainting, super-resolution and compressive sampling.

1 Introduction

State of the art image denoising methods perform a non-linear filtering that is adaptive to the image content. This adaptivity enables a non-local averaging of image features, thus making use of the relevant information along an edge or a regular texture pattern. This article shows how such adaptive non-local filterings can be extended to handle general inverse problems beyond simple denoising.

Adaptive Non-Local Image Processing. Traditional image processing methods use local computation over a time-frequency or a multi-scale domain [?]. These algorithms use either fixed transforms such as local DCT or wavelets, or fixed regularization spaces such as Sobolev or bounded variations, to perform image restoration.

In order to better respect edges in images, several edge-aware filtering schemes have been proposed, among which Yaroslavsky's filter [?], the bilateral filter [?], Susan filter [?] and Beltrami flow [?]. The non-local means filter [?] goes one step further by averaging pixels that can be arbitrary far away, using a similarity measure based on distance between patches. This non-local averaging shares similarities with patch-based computer graphics synthesis [?,?].

These adaptive filtering methods can be related to adaptive decompositions in dictionaries of orthogonal bases. For instance the bandlets best basis decomposition [?] re-transform the wavelet coefficients of an image in order to better capture edges. The grouplet transform of Mallat [?] does a similar retransformation but makes use of an adaptive geometric flow that is well suited to capture oriented oscillating textures [?].

⁰ This work was partially supported by ANR grant SURF-NT05-2.45825.

Regularization and inverse problems. Non-local filtering can be understood as a quadratic regularization based on a non-local graph, as detailed for instance in the geometric diffusion framework of Coifman et al. [?], which has been applied to non-local image denoising by Szlam et al. [?]. Denoising using quadratic penalization on image graphs is studied by Guilboa and Osher for image restoration and segmentation [?].

These quadratic regularizations can be extended to non-smooth energies such as the total variation on graphs. This has been defined over the continuous domain by Gilboa et al. [?] and over the discrete domain by Zhou and Schölkopf [?]. Elmoataz et al. [?] consider a larger class of non-smooth energies involving a p -laplacian for $p < 2$. Peyré replaces these non-linear flows on graphs by a non-iterative thresholding in a non-local spectral basis [?].

A difficult problem is to extend these graph-based regularizations to solve general inverse problems. The difficulty is that graph-based regularizations are adaptive since the graph depends on the image. To perform denoising, this graph can be directly estimated from the noisy image. To solve some specific inverse problems, the graph can also be estimated from the measurements. For instance, Kindermann et al. [?] and Buades et al. [?] perform image deblurring by using a non-local energy computed from the blurry observation. A similar strategy is used by Buades et al. [?] to perform demosaicing, where the non-local graph is estimated using an image with missing pixels. For inpainting of thin holes, Gilboa and Osher [?] compute a non-local graph directly from the image with missing pixels.

These approaches are different from recent exemplar-based methods introduced to solve some inverse problems, such as super-resolution (see for instance [?,?,?]). Although these methods operate on patches in a manner similar to non-local methods, they make use of pairs of exemplar patches where one knows both the low and high resolution version.

Contributions. This paper proposes a new framework to solve general inverse problems using a non-local and non-linear regularization on graphs. Our algorithm is able to efficiently solve for a minimizer of the proposed energy by iteratively computing an adapted graph and a solution of the inverse problem. We show applications to inpainting, super-resolution and compressive sampling where this new framework improves over wavelet and total variation regularization.

2 Non-local Regularization

2.1 Inverse Problems and Regularization

Many image processing problems can be formalized as the recovery of an image $f \in \mathbb{R}^n$ from a set of $p \leq n$ noisy linear measurements

$$y = \Phi f + \varepsilon \in \mathbb{R}^p.$$

where ε is an additive noise. The linear operator Φ typically accounts for some blurring, sub-sampling or missing pixels so that the measured data y only captures a small portion of the original image f one wishes to recover.

In order to solve this ill-posed problem, one needs to have some prior knowledge on the kind of typical images one expects to restore. This prior information should help to recover the missing information. Regularization theory assumes that f has some smoothness, for instance small derivatives (linear Sobolev regularization) or bounded variations (non-linear regularization).

A regularized solution f^* to the inverse problem can be written in variational form as

$$f^* = \operatorname{argmin}_{g \in \mathbb{R}^n} \frac{1}{2} \|y - \Phi g\|^2 + \lambda J(g), \tag{1}$$

where J is small when g is close to the smoothness model. The weight λ needs to be adapted to match the amplitude of the noise ε , which might be a non-trivial task in practical situations.

Classical variational priors include

- *Total variation*: The bounded variations model imposes that f^* has a small total variation and uses

$$J^{\text{tv}}(g) \stackrel{\text{def.}}{=} \|g\|_{\text{TV}} \stackrel{\text{def.}}{=} \int |\nabla_x g| dx. \tag{2}$$

This prior has been introduced by Rudin, Osher and Fatemi [?] for denoising purpose. It has been extended to solve many inverse problems, see for instance [?].

- *Sparsity priors*: Given a frame $(\psi_m)_m$ of \mathbb{R}^n , one defines a sparsity enforcing prior in this frame as

$$J^{\text{spars}}(g) \stackrel{\text{def.}}{=} \sum_m |\langle g, \psi_m \rangle|. \tag{3}$$

This prior has been introduced by Donoho and Johnstone [?] with the orthogonal wavelet basis for denoising purpose. It has then been used to solve more general inverse problems, see for instance [?] and the references therein. It can also be used in conjunction with redundant frames instead of orthogonal bases, see for instance [?,?].

2.2 Graph-based Regularization

Differential operators over graphs. We consider a weighted graph w that links together pixels x, y over the image domain with a weight $w(x, y)$. This graph allows to compute generalized discrete derivatives using the graph gradient operator

$$\forall x, \quad \nabla_x^w f = \left(\sqrt{w(x, y)} (f(y) - f(x)) \right)_y \in \mathbb{R}^n.$$

This operator defines, for any pixel x , a gradient vector $\nabla_x^w f \in \mathbb{R}^n$, see [?]. The divergence operator $\operatorname{div}^w = (\nabla^w)^T$ is the adjoint of the gradient, viewed as an

operator $f \mapsto \nabla^w f \in \mathbb{R}^{n \times n}$. For a gradient field $F_x \in \mathbb{R}^n$, the divergence is

$$(\operatorname{div}^w(F))(x) = \sum_y \sqrt{w(x,y)}(F_x(y) - F_y(x)).$$

The total-variation energy of an image, according to the graph structure given by w is then defined as

$$J_w(f) = \sum_x \|\nabla_x^w f\|, \quad (4)$$

where $\|\cdot\|$ is the euclidean norm over \mathbb{R}^n .

This energy was proposed by Gilboa et al. [?] in the continuous setting. It is used in the discrete setting by Zhou and Schölkopf [?] and Elmoataz et al. [?] in order to perform denoising.

Non-local graph adaptation. Given an image $f \in \mathbb{R}^n$ to process, one wishes to compute an adapted graph $w(f)$ so that the regularization by J_w efficiently removes noise without destroying the salient features of the image. In order to do so, we use a non-local graph inspired by the non-local means filtering [?], which has been used in several recent methods for denoising, see for instance [?,?]

This non-local graph is built by comparing patches around each pixel. A patch $p_x(f)$ of size $\tau \times \tau$ (τ being an odd integer) around a pixel position $x \in \{0, \dots, \sqrt{n} - 1\}^2$ is

$$\forall t \in \{-(\tau - 1)/2 + 1, \dots, (\tau - 1)/2\}^2, \quad p_x(f)(t) \stackrel{\text{def.}}{=} f(x + t).$$

A patch $p_x(f)$ is handled as a vector of size τ^2 . Color images f of n pixels can be handled using patches of dimension $3\tau^2$.

The non-local means algorithm [?] filters an image f using the following image-adapted weights $w = w(f)$

$$w(x, y) = \frac{\tilde{w}(x, y)}{Z_x} \quad \text{where} \quad \tilde{w}(x, y) = \begin{cases} e^{-\frac{\|p_x(f) - p_y(f)\|}{2\sigma^2}} & \text{if } \|x - y\| \leq \frac{\delta}{2}, \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where the normalizing constant is $Z_x \stackrel{\text{def.}}{=} \sum_y w(x, y)$. The parameter $\delta > 0$ restrict the non-locality of the method and also allows to speed-up computation. The parameter σ controls how many patch are taken into account to perform the averaging. It is a difficult parameter to set and ideally it should also be adapted to the noise level $|\varepsilon|$.

In the following, we consider the mapping $f \mapsto w(f)$ as a simple way to adapt a set of non-local weights to the image f to process.

Graph-based regularization of inverse problems. We propose to use this graph total-variation (??) to solve not only the denoising problem but arbitrary inverse problems such as inpainting, super-resolution and compressive sampling.

Our non-local graph regularization framework tries to recover an image f^* from a set of noisy measurements $y = \Phi f + \varepsilon$ using an adapted energy $J_{w(f)}$. The graph $w(f)$ should be adapted to the image f to recover, but unfortunately, one does not have this information since only the noisy observations y are available. In order to cope with such a problem, one performs an optimization over both the image f^* to recover and the optimal graph $w(f^*)$ as follow

$$f^* = \operatorname{argmin}_{g \in \mathbb{R}^n} \frac{1}{2} \|y - \Phi g\|^2 + \lambda J_{w(g)}(g). \quad (6)$$

It is important to note that the functional prior $J_{w(g)}$ depends non-linearly on the image g being recovered through equation (??).

2.3 Proximal Resolution of the Graph Regularization

The optimization of (??) is difficult because the energy $J_{w(g)}(g)$ makes it non-convex. We use an iterative approximate minimization that optimizes successively on the optimal graph and then on the image to recover.

Since J_w is a non-smooth functional, gradient descent methods are inefficient to minimize J_w . In order to cope with this issue, we use proximal iterations. The resulting algorithm is based on three main building blocks:

- The non-local graph adaptation procedure, equation (??), to compute a graph $w(f_k)$ adapted to the current estimate f_k of the algorithm.
- Proximal iterations to solve for (??) when the graph is fixed.
- Fixed point iterations to compute the proximity operator needed for the proximal iterations.

Proximal iterations. Equation (??) allows to compute a graph w adapted to a current estimate of the algorithm. In order to solve the initial regularization (??), we suppose that the graph w is fixed, and look for a minimizer of

$$f^*(w) = \operatorname{argmin}_{g \in \mathbb{R}^n} \frac{1}{2} \|y - \Phi g\|^2 + \lambda J_w(g). \quad (7)$$

This non-smooth convex minimization is difficult to solve because of the rank-deficient matrix Φ that couples the entries of g . In order to perform the optimization, we use iterative projections with a proximal operator.

Proximal iterations replace problem (??) by a series of simpler problems. This strategy has been developed to solve general convex optimizations and has been applied to solve non-smooth convex problems in image processing [?,?].

The proximity operator of a convex functional $J : \mathbb{R}^n \rightarrow \mathbb{R}^+$ is defined as

$$\operatorname{Prox}_J(f) = \operatorname{argmin}_{g \in \mathbb{R}^n} \frac{1}{2} \|f - g\|^2 + J(g). \quad (8)$$

A proximal iteration step uses the proximity operator to decrease the functional (??) one wishes to minimize

$$f^{(k+1)} = \operatorname{Prox}_{\frac{\lambda}{\mu} J_w} \left(f^{(k)} + \frac{1}{\mu} \Phi^T (y - \Phi f^{(k)}) \right). \quad (9)$$

It uses a step size $\mu > 0$ that should be set in order to ensure that $\|\Phi^T \Phi\| < \mu$. If the functional J is lipshitz continuous (which is the case of J_w), then one has the convergence $f^{(k)} \rightarrow f^*(w)$ a minimizer of (??), see for instance [?].

Computation of the proximity operator. In order to compute the proximity operator Prox_{J_w} for the functional J_w , one needs to solve the minimization (??), which is simpler than the original problem (??) since it does not involve anymore the operator Φ .

In order to do so, we use a fixed point algorithm similar to the one of Chambolle [?]. It is based on the computation of a gradient field $g(x) \in \mathbb{R}^n$ at each pixels x

$$g_{i+1}(x) = \frac{g_i(x) - \eta h(x)}{1 + \eta \|h(x)\|} \quad \text{with} \quad h(x) = \nabla_x^w \left(\text{div}^w(g_i) - \frac{\mu}{\lambda} f \right). \quad (10)$$

One can then prove that if η is small enough, the iterations (??) satisfy

$$f - \lambda \text{div}^w(g_i) \xrightarrow{i \rightarrow +\infty} \text{Prox}_{\frac{\lambda}{\mu} J_w}(f).$$

In numerical computation, since the graph w defined by equation (??) is relatively sparse (each pixel is connected to less than δ^2 pixels), $g(x)$ is stored as a sparse vector.

Graph regularization algorithm. The algorithm to minimize approximately (??) is detailed in table ???. It proceeds by iterating the proximal mapping (??). Each computation of this proximity operator requires m inner iterations of the dual gradient descent (??). Since the proximal iterations are robust against imperfect computation of the proximity operator, m is set to a small constant.

<ol style="list-style-type: none"> 1. <i>Initialization:</i> set $f^{(0)} = 0$ and $k \leftarrow 0$. 2. <i>Enforcing the constraints:</i> compute $\tilde{f}^{(k)} = f^{(k)} + \frac{1}{\mu} \Phi^T (y - \Phi f^{(k)})$. 3. <i>Update the graph:</i> compute the non-local graph $w^{(k)} = w(\tilde{f}^{(k)})$ adapted to $\tilde{f}^{(k)}$ using equation (??). 4. <i>Compute proximal iterations:</i> Set $g_0^{(k)} = \nabla^{w^{(k)}} f^{(k)}$. Perform m steps of the iterations of (??) $g_{i+1}^{(k)}(x) = \frac{g_i^{(k)}(x) - \eta h(x)}{1 + \eta \ h(x)\ } \quad \text{with} \quad h(x) = \nabla_x^{w^{(k)}} \left(\text{div}^{w^{(k)}}(g_i^{(k)}) - \frac{\mu}{\lambda} \tilde{f}^{(k)} \right).$ <p>Set the new estimate $f^{(k+1)} = \tilde{f}^{(k)} - \lambda \text{div}^{w^{(k)}}(g_m^{(k)})$.</p> 5. <i>Stopping criterion:</i> while not converged, set $k \leftarrow k + 1$ and go back to 2.

Table 1: Block coordinate descent algorithm to minimize approximately (??).

3 Numerical Illustration

In the numerical simulations, we consider three different regularizations:

- The total variation energy J^{tv} , defined in equation (??). An algorithm very similar to the algorithm of table ?? is used for this minimization, excepted that ∇_x is the classical gradient and that the step 3 of algorithm ?? is not needed.
- The sparsity energy J^{spars} , defined in equation (??), using a redundant tight frame of translation invariant wavelets $(\psi_m)_m$. An algorithm very similar to the one of table ?? is used for this minimization, excepted that the proximal projection is computed with a soft thresholding as detailed in [?] and that the step 2 of algorithm ?? is not needed.
- The non-local total variation regularization J_w in an optimized graph, solved using algorithm ?. For this regularization, the parameter σ of equation (??) is fixed by hand in order to have consistent results for all experiments. The locality parameter δ of equation (??) is fixed to 15 pixels.

Both total variation and non-local total variation require approximately the same number of iterations. For these two methods, the number of inner iterations to solve for the proximity operator is set to $m = 10$. The non-local iteration is computationally more intensive since the computation of the non-local weights $(w(x, y))_y$ requires to explore δ^2 pixels y for each pixel x . In the three applications of sections ??, ?? and ??, we use a low noise level $|\varepsilon|$ of $.02\|y\|$. For all the proposed methods, the parameter λ is optimized in an oracle manner in order to minimize the PSNR of the recovered image f^*

$$\text{PSNR}(f^*, f) = -20 \log_2(\|f^* - f\|/\|f\|_\infty).$$

3.1 Inpainting

Inpainting corresponds to the operation of removing pixels from an image

$$(\Phi f)(x) = \begin{cases} 0 & \text{if } x \in \Omega, \\ f(x) & \text{if } x \notin \Omega, \end{cases}$$

where $\Omega \subset \{0, \dots, \sqrt{n}-1\}^2$ is the region where the input data has been damaged. In this case, $\Phi^T = \Phi$, and one can take a proximity step size $\mu = 1$ so that the proximal iteration (??) becomes

$$f^{(k+1)} = \text{Prox}_{\lambda J}(\tilde{f}^{(k)}) \quad \text{where} \quad \tilde{f}^{(k)}(x) = \begin{cases} f(x) & \text{if } x \in \Omega, \\ f^{(k)}(x) & \text{if } x \notin \Omega. \end{cases}$$

Classical methods for inpainting use partial differential equations that propagate the information from the boundary of Ω to its interior, see for instance [?,?,?,?]. Sparsity promoting prior such as (??) in wavelets frames and local cosine bases have been used to solve inpainting as a inverse problem [?,?].

Our method iteratively updates the optimal graph $w(x, y)$ and iterations of the algorithm update in parallel all the pixels inside Ω . This is similar to

the exemplar-based inpainting algorithm of [?] which uses a non-local copying of patches, however in their framework Ω is filled progressively by propagating inward from the boundary. The anisotropic diffusion of Tschumperle and Deriche [?] also progressively builds an adapted operator (parameterized by a tensor field) but they solve a PDE and not a regularization as we do.

Figure ?? shows some numerical examples of inpainting on images where 80% of the pixels have been damaged. The wavelets method performs better than total variation in term of PSNR but tends to introduce some ringing artifact. The non-local total variation perform better in term of PSNR and is visually more pleasing since edge are better reconstructed.

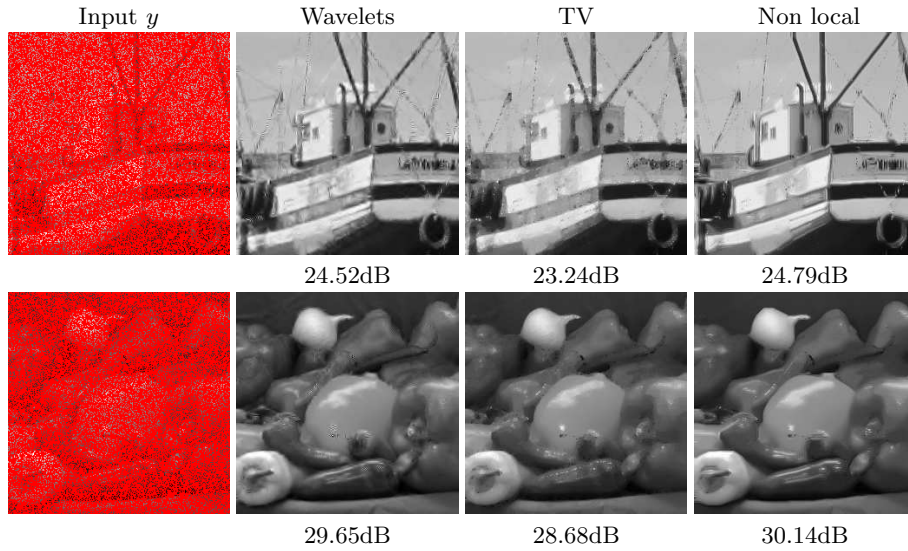


Fig. 1. Examples of inpainting where Ω occupies 80% of pixels. The original images f are displayed on the left of figure ??.

3.2 Super-resolution

Super-resolution corresponds to the recovery of a high-definition image from a filtered and sub-sampled image. It is usually applied to a sequence of images in video, see the review papers [?,?]. We consider here a simpler problem of increasing the resolution of a single still image, which corresponds to the inversion of the operator

$$\forall f \in \mathbb{R}^n, \quad \Phi f = (f * h) \downarrow^k \quad \text{and} \quad \forall g \in \mathbb{R}^p, \quad \Phi^T g = (g \uparrow^k) * h$$

where $p = n/k^2$, $h \in \mathbb{R}^n$ is a low-pass filter, $\downarrow^k: \mathbb{R}^n \rightarrow \mathbb{R}^p$ is the sub-sampling operator by a factor k along each axis and $\uparrow^k: \mathbb{R}^p \rightarrow \mathbb{R}^n$ corresponds to the insertion of $k - 1$ zeros along horizontal and vertical directions.




Fig. 2. Examples of image super-resolution with a down-sampling $k = 8$. The original images f are displayed on the left of figure ??.

Figure ?? shows some graphical results of the three tested super-resolution methods. The results are similar to those of inpainting, since our method improves over both wavelets and total variation.

3.3 Compressive-sampling

Compressive sensing is a new sampling theory that uses a fixed set of linear measurements together with a non-linear reconstruction [?,?]. The sensing operator computes the projection of the data on a finite set of p vectors

$$\Phi f = \{ \langle f, \varphi_i \rangle \}_{i=0}^{p-1} \in \mathbb{R}^p, \tag{11}$$

where $(\varphi_i)_{i=0}^{p-1}$ are the rows of Φ .

Compressive sampling theory gives hypotheses on both the input signal f and the sensing vectors $(\varphi_i)_i$ for this non-uniform sampling process to be invertible. In particular, the $(\varphi_i)_i$ must be incoherent with the orthogonal basis $(\psi_m)_m$ used for the sparsity prior, which is the case with high probability if they are drawn randomly from unit normed random vectors. Under the additional condition that f is sparse in an orthogonal basis $(\psi_m)_m \# \{m \mid \langle f, \psi_m \rangle \neq 0\} \leq s$ then

the optimization of (??) using the energy (??) leads to a recovery with a small error $\|\bar{f} - f\| \approx |\varepsilon|$ if $p = O(s \log(n/s))$. This results extend to approximately sparse signals, such as for instance signals that are highly compressible in an orthogonal basis.

In the numerical tests, we choose the columns of $\Phi \in \mathbb{R}^{p \times m}$ to be independent random unit normed vectors. Figure ?? shows examples of compressive sampling reconstructions. The results are slightly above the wavelets method and tend to be visually more pleasing.




Fig. 3. *Examples of compressed sensing reconstruction with $p = n/8$.*

Conclusion and Future Work

This paper proposed a new framework for the non-local resolution of linear inverse problems. The variational minimization computes iteratively an adaptive non-local graph that enhances the geometric features of the recovered image. Numerical tests show how this method improves over some state of the art methods for inpainting, super-resolution and compressive sampling. This new method also open interesting questions concerning the optimization of the non-local graph. While this paper proposes to adapt the graph using a patch comparison principle, it is important to understand how this adaptation can be re-casted as a variational minimization.