

A two-step mixed inpainting method with curvature-based anisotropy and spatial adaptivity

F. J. IBARROLA

Instituto de Investigación en Señales, Sistemas e Inteligencia Computacional, sinc(i),
FICH-UNL/CONICET, Argentina.
Ciudad Universitaria, CC 217, Ruta Nac. No 168, km 472.4, (3000) Santa Fe, Argentina.
fibarrola@sinc.unl.edu.ar

R. D. SPIES

Instituto de Matemática Aplicada del Litoral, IMAL, CONICET-UNL, Centro Científico
Tecnológico CONICET Santa Fe, Colectora Ruta Nac. 168, km 472, Paraje “El Pozo”,
3000, Santa Fe, Argentina and Departamento de Matemática, Facultad de Ingeniería Química,
Universidad Nacional del Litoral, Santa Fe, Argentina. rspies@santafe-conicet.gov.ar

Abstract

The image inpainting problem consists of restoring an image from a (possibly noisy) observation, in which data from one or more regions are missing. Several inpainting models to perform this task have been developed, and although some of them perform reasonably well in certain types of images, quite a few issues are yet to be sorted out. For instance, if the image is expected to be smooth, the inpainting can be made with very good results by means of a Bayesian approach and a maximum a posteriori computation [2]. For non-smooth images, however, such an approach is far from being satisfactory. Even though the introduction of anisotropy by prior smooth gradient inpainting to the latter methodology is known to produce satisfactory results for slim missing regions [2], the quality of the restoration decays as the occluded regions widen. On the other hand, Total Variation (TV) inpainting models based on high order PDE diffusion equations can be used whenever edge restoration is a priority. More recently, the introduction of spatially variant conductivity coefficients on these models, such as in the case of Curvature-Driven Diffusion (CDD) [4], has allowed inpainted images with well defined edges and enhanced object connectivity. The CDD approach, nonetheless, is not quite suitable wherever the image is smooth, as it tends to produce piecewise constant restorations.

In this work we present a two-step inpainting process. The first step consists of using a CDD inpainting to build a pilot image from which to infer *a-priori* structural information on the image gradient. The second step is inpainting the image by minimizing a mixed spatially variant anisotropic functional, whose weight and penalization directions are based upon the aforementioned pilot image. Results are presented along with comparison measures in order to illustrate the performance of this inpainting method.

keywords: Inpainting, Inverse Problems, Ill-posedness, Regularization, Anisotropy.

⁰The final version of this article was published in *Inverse Problems and Imaging* Volume 11, No. 2, 2017, 247262

1 INTRODUCTION

Let us consider an image from which data from one or more regions are missing. The inpainting problem can be described as that of using the available information from the occluded image for filling up the missing areas in a way that looks natural and intuitive to the human eye. Although it is clear then that defining a “good” inpainting method involves a high degree of subjectivity, some properties are to be expected. For instance, if an occluded region runs over a smooth area, we would expect the inpainting to be smooth, while if it goes through an edge we would expect the two parts set apart by the occlusion to result “appropriately” connected. These desired properties are often hard to include in a mathematical model, even more so, simultaneously. In this article we shall combine the Curvature Driven Diffusion (CDD) model proposed in [4] with the ideas of including structural prior information [2] into regularization methods to develop a two-step inpainting process which adequately complies with the aforementioned criteria.

We shall begin by fixing some notation. Our (grayscale) image will be associated to a function $u : \Omega \subset \mathbb{R}^2 \rightarrow [0, 1]$, where $u(x, y)$ represents the light intensity of the point (x, y) , corresponding $u = 1$ to white, and $u = 0$ to black. With D we shall denote the occlusion and with $v \doteq u|_{\Omega \setminus D}$ the known part of u .

Before we proceed to build our two-step method, it is timely to describe three traditional inpainting methods that are closely bonded to it.

1.1 Tikhonov-Phillips Inpainting

If the image to be inpainted is expected to be smooth, then there are traditional methods that turn out to be quite satisfactory. With the Tikhonov-Phillips regularization method of order 1 (T1), the inpainted image is defined as

$$\hat{u} \doteq \arg \min_{u \in L^2(\Omega)} \{ \|\mathcal{T}u - v\|_{L^2(\Omega \setminus D)}^2 + \lambda \|\nabla u\|_{L^2(\Omega)}^2 \}, \quad (1)$$

where $\mathcal{T} : L^2(\Omega) \rightarrow L^2(\Omega \setminus D)$ is the D -occlusion operator, defined by $\mathcal{T}u = u|_{\Omega \setminus D}$, and $\lambda > 0$ is a regularization parameter. It is well known, however, that if the original image has edges or borders set apart by an occlusion, then such an inpainting method fails to extend them inside the occluded region (see Figure 1(b)). Edge preservation, however, can be fairly enhanced by appropriately modifying the penalizing term in (1). One way of doing so is to introduce an anisotropy-inducing matrix field A [2], which attenuates penalization on the maximum gradient direction. In this case, the inpainted image is defined as

$$\hat{u} \doteq \arg \min_{u \in L^2(\Omega)} \{ \|\mathcal{T}u - v\|_{L^2(\Omega \setminus D)}^2 + \lambda \|A \nabla u\|_{L^2(\Omega)}^2 \}. \quad (2)$$

The matrix field A makes the penalizer take into account not only the norm of the gradient, but also its direction. The properties of A and a way to build it will be presented later, but it is timely to point out here that its construction requires *a-priori* information of the expected image. Figure 1(c) depicts an order 1 Tikhonov-Phillips anisotropic inpainting for which A was build upon an *a-priori* order 2 Tikhonov-Phillips inpainting of the gradient field, as suggested in [2]. Although we can see that this method works significantly better than the isotropic one for “thin” occluded regions, it becomes inappropriate as those regions widen (see Figure 1(c)).

We then describe an inpainting method known to possess better edge preserving properties.




Figure 1: (a) Occluded image; (b) T1 inpainting; (c) Anisotropic T1 inpainting.

1.2 Total Variation inpainting

One of the most commonly used edge preserving methods is the so called total variation (TV) regularization [1], [10]. With this method, the inpainted image is given by

$$\hat{u} \doteq \arg \min_{u \in L^2(\Omega)} \{ \|\mathcal{T}u - v\|_{L^2(\Omega \setminus D)}^2 + \lambda \|\nabla u\|_{L^1(\Omega)} \}. \quad (3)$$

By using the corresponding Euler-Lagrange equation, the minimizer in (3) can be thought of as the steady-state solution of the following diffusion PDE (see [7])

$$\frac{\partial u}{\partial t} = \nabla \cdot \left[\frac{\nabla u}{|\nabla u|} \right] + \frac{2}{\lambda} (\mathcal{T}u - v). \quad (4)$$




Figure 2: Occluded image (a); TV inpainting (b).

TV inpainting tends to produce piecewise constant solutions with sharp edges inside the occlusion (this can be seen in Figure 2). Even though this inpainting

process seems to fairly comply with the edge preservation requirement, it presents a peculiar property. In fact, if we look at the black vertical bar in Figure 2(a), we notice that it is wider than the occluded region, but whenever the occluded region is wider than the bar (Figure 3(a)), the TV solution is very different (see Figure 3(b)). Although there is no formal reason to sustain that the inpainted image in Figure 3(b) is not a “good” solution, it is a fact that most humans will prefer the image in Figure 2(b) over the one in Figure 3(b) as an inpainting of the image in Figure 3(a). This property of favoring the linking of similar objects on opposite sides of an occlusion is referred to as *object connectivity*.




Figure 3: (a) Occluded image; (b) TV inpainting.

Once again, this puts into evidence that any inpainting process can entail a high level of subjectivity. That being said, a “good” inpainting method could be conceived as one that would most frequently emulate what most humans would do. The rest of the article is strongly aligned with this belief.

1.3 Curvature Driven Diffusion Inpainting

An *isophote* of an image represented by u can be roughly thought of as any level line of u that separates regions of different light intensities. It is highly desired for an inpainting method to be able to connect isophotes inside the occluded areas, since this will result in that the edges “broken” by the occlusion will be reconnected. In fact, in Figure 3(b) we observe that the original isophotes (the edges of the black bar) do not result connected as we would hope. Note also that the isophotes of the inpainted image have corners, meaning that their curvature κ at those points is $\pm\infty$, in contrast with the zero curvature in the isophotes of the “expected” inpainted image (Figure 2(b)). This observation led Chan and Shen [4] to include the curvature into the diffusion model (4). By letting $\hat{D} = \frac{1}{|\nabla u|}$, equation (4) restricted to the occlusion reads

$$\frac{\partial u}{\partial t} = \nabla \cdot [\hat{D} \nabla u].$$

The ideas in [4] consist essentially in appropriately modifying the diffusion coefficient \hat{D} in order to take curvature into account. This can be done by redefining $\hat{D} = \frac{g(|\kappa|)}{|\nabla u|}$,

where $g : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ is an increasing function such that $g(0) = 0$ and $g(\infty) = \infty$. In this way, diffusion is strong where curvature is large, while it is weak where curvature is small. This modification leads to the so called Curvature-Driven Diffusion (CDD) equation

$$\frac{\partial u}{\partial t} = \nabla \cdot \left[\frac{g(|\kappa|)}{|\nabla u|} \nabla u \right]. \quad (5)$$

Now, since the curvature of the level lines is $\pm\infty$ at the corners of the TV inpainted image in Figure 3(b), clearly such an image cannot be the steady-state of equation (5). It turns out that in this case the method strongly favors connectivity on the steady state, as it can be seen in Figure 4. Nonetheless, it occurs that this method tends to produce piecewise constant restorations, and therefore it results unsuitable for inpainting over smooth regions (see Figure 4).




Figure 4: (a) Occluded image; (b) CDD inpainting.

Before we go any further, we should mention that the presence of noise is common in this kind of problems. The CDD inpainting model (5) can take noise into account, for instance, by using TV regularization outside the occluded region. Thus, the resulting diffusion PDE takes the form

$$\frac{\partial u}{\partial t} = \nabla \cdot \left[\frac{g(|\kappa|)\chi_D + \chi_{\Omega \setminus D}}{|\nabla u|} \nabla u \right] + \frac{2}{\lambda}(u - v)\chi_{\Omega \setminus D}. \quad (6)$$

Note that (6) reduces to equation (5) inside the occlusion, and to (4) outside of it.

The main objective of this article is to develop an inpainting model that complies with the three aforementioned properties. Namely, it must preserve edges while enhancing object connectivity, and allowing for smooth inpainting wherever it is deemed appropriate.

2 INPAINTING VIA COMBINED CDD AND T1-TV METHODS

Regularization methods can be combined by using spatially-varying weighted averages of two or more penalizers [8] to encompass the characteristic properties of

each one. We shall consider the case of mixed T1 and TV regularization, taking advantage of the fact that the first one tends to produce smooth restorations, while the latter is better suited for edge preservation. The inpainted image will then be defined as the minimizer of

$$\mathcal{J}(u) \doteq \|\mathcal{T}u - v\|_{L^2(\Omega \setminus D)}^2 + \lambda_{T1} \|\sqrt{1 - \theta} A \nabla u\|_{L^2(\Omega)}^2 + \lambda_{TV} \|\theta A \nabla u\|_{L^1(\Omega)}, \quad (7)$$

where $\lambda_{T1} > 0$ and $\lambda_{TV} > 0$ are regularization parameters, $A = A(x, y)$ is an anisotropy matrix field and $\theta = \theta(x, y) \in [0, 1]$ is a spatially-varying function weighting both penalizers at each point. Note that $\theta = 0$ corresponds to pure anisotropic T1 regularization, while $\theta = 1$ leads to pure anisotropic TV regularization.

The matrix field $A : \Omega \rightarrow \mathbb{R}^{2 \times 2}$ is used for introducing anisotropy to the regularization model. One way of constructing A can be found in [2], where it is built upon the gradient field of an *a-priori* estimation $u_p(x, y)$ of $u(x, y)$. The first step for doing this is to define a continuous and decreasing function $h : \mathbb{R}_0^+ \rightarrow (0, 1]$, satisfying $h(0) = 1$ and $\lim_{t \rightarrow \infty} h(t) = 0$, that will serve to determine the eigenvalues of $A(x, y)$, constructed as follows:

$$A(x, y) = I - (1 - h(|\nabla u_p(x, y)|)) \begin{bmatrix} \nabla u_p(x, y) \\ |\nabla u_p(x, y)| \end{bmatrix} \begin{bmatrix} \nabla u_p(x, y) \\ |\nabla u_p(x, y)| \end{bmatrix}^T. \quad (8)$$

As a consequence, A has the following properties:

- $A(x, y)$ is a symmetric positive definite matrix $\forall (x, y) \in \Omega$.
- If $\nabla u_p(x, y) = 0$, $A(x, y) = I$ (the identity matrix).
- If $\nabla u_p(x, y) \neq 0$, $A(x, y)$ has eigenvalues $\sigma_j(x, y)$ and eigenvectors $v_j(x, y)$, $j = 1, 2$, such that

$$\begin{aligned} v_1(x, y) &\parallel \nabla u_p(x, y), & \sigma_1(x, y) &= h(|\nabla u_p(x, y)|), \\ v_2(x, y) &\perp \nabla u_p(x, y), & \sigma_2(x, y) &= 1. \end{aligned}$$

With the required properties on h , the eigenvalues $\sigma_1(x, y)$ turn out to be small at points (x, y) where the norm of the estimated gradient $\nabla u_p(x, y)$ is large, while $\sigma_2(x, y)$ remains constant. This is desirable since it will later translate into a decay on penalization in the expected gradient direction while keeping it unchanged in its normal direction. The function h can be chosen, for instance, as $h(t) = 1/(1 + (t/\tau)^k)$, where $\tau, k > 0$ are control parameters that could be roughly thought of as a lower threshold for the values of $|\nabla u_p|$ starting from which we infer the image has an edge and the width of the transition region, respectively.

In [2], the gradient field estimation ∇u_p was built by inpainting the components of the gradient using a second order Tikhonov-Phillips method. However, since this method produces smooth restorations, edge preservation remains compromised for wide occluded regions, as it can be seen in Figure 1(c). To overcome this disadvantage, we shall build the matrix field A based upon the gradients of an *a-priori* CDD inpainting of the image, which as previously noted, favors object connectivity.

In the same fashion, the weighting function θ , which also requires of *a-priori* information about the local regularity of u , can be built from a ‘‘pilot image’’ u_p ,

obtained by means of a CDD-inpainting process. We take $\theta(x, y) = w(|\nabla u_p(x, y)|)$, where w is an appropriately chosen increasing function, with $0 \leq w(t) \leq 1 \forall t \in \mathbb{R}_0^+$.

Due to the way in which the pilot image is built, if used directly as obtained by the CDD-inpainting process, the staircasing effect on u_p will have a negative impact on the effect of the weighting and anisotropy functions, which in turn might lead to sharp artificial edges appearing in the final inpainted image. To overcome this, the pilot image is smoothed out with a small-variance Gaussian kernel.

Our full inpainting process can then be stated as follows:

Step 1: CDD inpainting. Perform a CDD inpainting on the occluded image to obtain a first estimation. That is, compute \tilde{u}_p as the steady state of the equation

$$\frac{\partial u}{\partial t} = \nabla \cdot \left[\frac{g(|\kappa|)\chi_D + \chi_{\Omega \setminus D}}{|\nabla u|} \nabla u \right] + \frac{2}{\lambda}(u - v)\chi_{\Omega \setminus D}.$$

Step 2: Pilot image smoothing. Smooth out \tilde{u}_p by means of a low-pass filter by computing

$$u_p \doteq G * \tilde{u}_p,$$

where G is a low-variance Gaussian kernel, and “ $*$ ” denotes convolution. This step will result in the good definition of the gradient field ∇u_p , while attenuating the staircasing effect that may appear in the final restoration.

Step 3: Construction of the anisotropy matrix field. Use u_p to construct A as

$$A = I - \left(1 - \frac{1}{1 + (|\nabla u_p|/\tau)^k} \right) \left[\frac{\nabla u_p}{|\nabla u_p|} \right] \left[\frac{\nabla u_p}{|\nabla u_p|} \right]^T.$$

Step 4: Construction of the weighting function. Use u_p to build θ as

$$\theta(x, y) = \frac{|\nabla u_p(x, y)|}{\max_{(w, z) \in \Omega} |\nabla u_p(w, z)|}.$$

Step 5: Final inpainting. Use the previously computed A and θ to build the mixed weighted T1-TV anisotropic functional and compute the restored image as

$$\hat{u} = \arg \min_{u \in L^2(\Omega)} \{ \|Tu - v\|_{L^2}^2 + \lambda_{T1} \| \sqrt{1 - \theta} A \nabla u \|_{L^2}^2 + \lambda_{TV} \| \theta A \nabla u \|_{L^1} \}. \quad (9)$$

Having stated our new inpainting method, we shall proceed to briefly describe an appropriate numerical implementation.

3 NUMERICAL IMPLEMENTATION

We start the implementation of our inpainting method by performing a discretization over the image domain. We assume that the grayscale image domain is $\Omega = [0, 1] \times [0, 1]$ and we discretize it to obtain an M -by- M pixel grid and an M -by- M matrix U , consisting its entries of the values of the function u at the centerpoints of the pixels. Next, we stack the columns of the matrix U to get a vector $u \in \mathbb{R}^{M^2}$ so that $u_{M(l-1)+m} = U_{m,l} \forall l, m = 1, 2, \dots, M$. For better understanding we will often identify $u_{M(l-1)+m}$ with $u(x, y)$.

3.1 The CDD method

Let us define $\xi = (\xi_1, \xi_2)^T \doteq \frac{g(|\kappa|)}{|\nabla u|} \nabla u$, whose divergence we need to approximate. We do it by computing

$$\nabla \cdot \xi(x, y) \approx \frac{\xi_1(x + h/2, y) - \xi_1(x - h/2, y)}{h} + \frac{\xi_2(x, y + h/2) - \xi_2(x, y - h/2)}{h}, \quad (10)$$

where $h = 1/M$ is the pixel-width.

Now, in order to compute (10) at the midpoints between adjacent pixels, we need to estimate both the gradient of u and the curvature κ at those points. Firstly, for the points of the form $\nabla u(x + h/2, y)$, we compute the components of ∇u as

$$\begin{aligned} u_x(x + \frac{h}{2}, y) &\approx \frac{u(x + h, y) - u(x, y)}{h}, \\ u_y(x + \frac{h}{2}, y) &\approx \frac{\frac{1}{2}(u(x + h, y + h) + u(x, y + h)) - \frac{1}{2}(u(x + h, y - h) + u(x, y - h))}{2h}. \end{aligned}$$

At the points of the form $(x, y + h/2)$ the construction is analogous.

As for the isophote curvature κ , we compute it explicitly as a function of the gradient of u :

$$\kappa = \nabla \cdot \left[\frac{\nabla u}{|\nabla u|} \right] = \frac{\partial}{\partial x} \frac{u_x}{|\nabla u|} + \frac{\partial}{\partial y} \frac{u_y}{|\nabla u|}.$$

At points of the form $(x + h/2, y)$, we approximate it with

$$\begin{aligned} \kappa(x + h/2, y) &= \frac{1}{h} \left[\frac{u_x(x + h, y)}{|\nabla u(x + h, y)|} - \frac{u_x(x, y)}{|\nabla u(x, y)|} \right] \\ &\quad + \frac{1}{2h} \left[\frac{u_y(x + h/2, y + h)}{|\nabla u(x + h/2, y + h)|} - \frac{u_y(x + h/2, y - h)}{|\nabla u(x + h/2, y - h)|} \right]. \end{aligned} \quad (11)$$

The computation of κ at points of the form $(x, y + h/2)$ is analogous. To avoid division by zero in the last equation (when $\nabla u = 0$), we replace $|s|$ by $\sqrt{s^2 + \epsilon^2}$, with $\epsilon > 0$ chosen sufficiently small.

We are then ready to state the algorithm for the CDD inpainting process. It is timely to point out here that due to numerical stability issues, instead of using Euler's method to solve the IVP (as stated in [4]), we use an Adams-Moulton Adams-Bashford predictor-corrector method, which showed better performance. Our modified iterative algorithm reads as follows:

Step 1: Initializing. Define an arbitrary initial estimation $u^{(0)}$, coinciding with the data outside the occlusion, and let $n = 0$.

Step 2: Updating. For $m : 1 \dots M$, define

$$f(\kappa_m^{(n)}, u_m^{(n)}) \doteq \nabla \cdot \left[\frac{g(|\kappa_m^{(n)}|) \chi_D + \chi_{\Omega \setminus D}}{|\nabla u_m^{(n)}|} \nabla u_m^{(n)} \right] + \frac{2}{\lambda} (u_m^{(n)} - v_m) \chi_{\Omega \setminus D}, \quad (12)$$

and compute $u_m^{(n+1)}$ as follows:

$$\begin{aligned} \nu_m^{(n+1)} &= u_m^{(n)} + \frac{\Delta t}{12} \left[23f(\kappa_m^{(n)}, u_m^{(n)}) - 16f(\kappa_m^{(n-1)}, u_m^{(n-1)}) + 5f(\kappa_m^{(n-2)}, u_m^{(n-2)}) \right] \\ u_m^{(n+1)} &= u_m^{(n)} + \frac{\Delta t}{12} \left[5f(\kappa_m^{(n+1)}, \nu_m^{(n+1)}) + 8f(\kappa_m^{(n)}, u_m^{(n)}) - f(\kappa_m^{(n-1)}, u_m^{(n-1)}) \right] \end{aligned}$$

where for each m , $\kappa^{(n)}$ and $\nabla u^{(n)}$ are computed from $u^{(n)}$ as previously described. A second order Runge-Kutta method is used for the first two steps.

Step 3: Stopping. If an appropriate stopping criterion (defined upon the decay of the sum of the curvature at each pixel) is reached, the process stops and the inpainted image is defined as $u^{(n+1)}$. Otherwise, n is increased ($n = n + 1$) and the algorithm continues from Step 2.

Next, we show how the mixed weighted anisotropic T1-TV regularization can be numerically implemented.

3.2 Mixed anisotropic T1-TV implementation

To find the minimizer of the T1-TV inpainting functional given in (7), we consider the discretized version

$$J(u) = \frac{1}{M^2} \|Tu - v\|^2 + \frac{\lambda_{T1}}{M^2} \sum_{m \in \mathfrak{M}} (1 - \theta_m) \left\| A_m \begin{pmatrix} M(u_m - u_{m+1}) \\ M(u_m - u_{m-M}) \end{pmatrix} \right\|_2^2 + \frac{\lambda_{TV}}{M^2} \sum_{m \in \mathfrak{M}} \theta_m \left\| A_m \begin{pmatrix} M(u_m - u_{m+1}) \\ M(u_m - u_{m-M}) \end{pmatrix} \right\|_2, \quad (13)$$

where $T \in \mathbb{R}^{M^2 \times M^2}$ is the diagonal matrix associated to the occlusion operator \mathcal{T} , θ_m and A_m are the evaluations of the weighting function and of the anisotropy matrix field at the centerpoint of the m^{th} pixel, and \mathfrak{M} denotes the set of indices corresponding to the interior pixels, on which the gradient is estimated, that is $\mathfrak{M} \doteq \{m \in \mathbb{N} : M < m < M(M + 1), Mk \neq m \neq Mk + 1 \forall k = 1 \dots M\}$.

Although there are several efficient ways to estimate the minimizer of (13), we have chosen a half-quadratic approach, which is described below, for it has proven to cope very well with the great dimensions that this problem can present (details can be found in [5]).

First, we approximate J by a differentiable functional and make use of a duality relation to iteratively approach its minimizer. We begin by replacing the last term in (13) by a differentiable approximation in order to find the minimizer by considering the first order necessary condition. We do so by replacing, for $w \in \mathbb{R}^2$, the value of $\|w\|_2$ by $\phi(\|w\|_2)$, where $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is given by $\phi(t) \doteq \sqrt{t^2 + \eta^2} - \eta$, for η sufficiently small. With this choice of ϕ , it can be shown ([9]) that there exists a function ψ satisfying the following duality relation

$$\begin{aligned} \phi(t) &= \inf_{s > 0} (st^2 + \psi(s)), \\ \psi(s) &= \sup_{t \in \mathbb{R}} (\phi(t) - st^2), \end{aligned} \quad (14)$$

and therefore

$$\left\| A_m \begin{pmatrix} M(u_m - u_{m+1}) \\ M(u_m - u_{m-M}) \end{pmatrix} \right\|_2 \approx \inf_{s_m \in \mathbb{R}^+} \{s_m (t_{m,1}^2 + t_{m,2}^2) + \psi(s_m)\}, \quad (15)$$

where

$$t_{m,1} = \frac{a_{1,1}^m (u_m - u_{m-M}) + a_{1,2}^m (u_m - u_{m+1})}{1/M}, \quad (16)$$

and

$$t_{m,2} = \frac{a_{2,1}^m(u_m - u_{m-M}) + a_{2,2}^m(u_m - u_{m+1})}{1/M}. \quad (17)$$

Define now the M^2 -by- M^2 diagonal matrices $A^{i,j}$, for $i, j = 1, 2$, such that $A_{m,m}^{i,j} = a_{i,j}^m$ if $m \in \mathfrak{M}$ and $A_{m,m}^{i,j} = 0$ otherwise. In a similar fashion, let $\Theta \doteq \text{diag}(\theta_m)_{M^2 \times M^2}$ and $S \doteq \text{diag}(s_m)_{M^2 \times M^2}$. Let L_x and L_y be the M^2 -by- M^2 first order finite difference approximating matrices for the components of the gradient, and let R_1 and R_2 be the M^2 -by- M^2 matrices defined as $R_1 \doteq A^{1,1}L_x + A^{1,2}L_y$ and $R_2 \doteq A^{2,1}L_x + A^{2,2}L_y$. Finally, let I be the M^2 -by- M^2 identity matrix, and define the functional

$$\begin{aligned} K_{\theta,\phi}(u, s) &\doteq \frac{1}{M^2} \|Tu - v\|^2 + \frac{\lambda_{T1}}{M^2} u^T (R_1^T (I - \Theta) R_1 + R_2^T (I - \Theta) R_2) u + \\ &\quad + \frac{\lambda_{TV}}{M^2} u^T (R_1^T \Theta S R_1 + R_2^T \Theta S R_2) u + \frac{\lambda_{TV}}{M^2} \sum \theta_m \psi(s_m). \end{aligned} \quad (18)$$

It can be shown ([5]) that

$$\inf_{s \in \mathbb{R}^{M^2}} K_{\theta,\phi}(u, s) = J_\phi(u) \approx J(u), \quad (19)$$

where J_ϕ is the functional obtained from J by using approximation (15). Hence, our problem turns out to be equivalent to minimizing K with respect to both u and s simultaneously. Note that the first order necessary condition on K with respect to u can be written as

$$\left(T^T T + \lambda_{T1} (R_1^T (I - \Theta) R_1 + R_2^T (I - \Theta) R_2) + \lambda_{TV} R_1^T \Theta S R_1 + \lambda_{TV} R_2^T \Theta S R_2 \right) u = T^T v. \quad (20)$$

In order to minimize $K_{\theta,\phi}(u, s)$ with respect to s , we define

$$b_m \doteq \arg \min_{s_m \in \mathbb{R}_+} \{ s_m (t_{m,1}^2 + t_{m,2}^2) + \psi(s_m) \},$$

where $t_{m,1}$ and $t_{m,2}$ are defined as in (16) and (17), respectively, and resort to (14) to deduce that b_m must satisfy

$$b_m = \frac{\phi' \left(\sqrt{t_{m,1}^2 + t_{m,2}^2} \right)}{2 \sqrt{t_{m,1}^2 + t_{m,2}^2}}. \quad (21)$$

Although details on the derivation of equation (21) can be found in [6], we give here a brief sketch of the proof. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined as $f(s, t) \doteq st^2 + \psi(s)$. Then, if $b \doteq \arg \min_{s \in \mathbb{R}_+} f(s, t)$, the first order necessary condition over f with respect to s yields $0 = t^2 + \frac{\partial \psi(b)}{\partial b}$. Notice that $\phi(t) = f(b, t)$, and hence

$$\frac{\partial \phi(t)}{\partial t} = \frac{\partial f(b, t)}{\partial t} = 2bt + \frac{\partial b}{\partial t} t^2 + \frac{\partial \psi(b)}{\partial b} \frac{\partial b}{\partial t} = 2bt + \frac{\partial b}{\partial t} \left(t^2 + \frac{\partial \psi(b)}{\partial b} \right) = 2bt.$$

Finally, we state our cyclic iterative algorithm for the T1-TV image inpainting as follows:

Step 1: Initializing. Set $j = 0$, and initialize $u^j = u^0$ (e.g. $u^0|_{\Omega \setminus D} = v$ and $u^0|_D = 0$) and $b^j = b^0$ (e.g. $b^0 = 0$).

Step 2: Counting. Make $j = j + 1$.

Step 3: Updating b . Update b^j using equation (21):

$$b_m^j = \frac{\phi' \left(\sqrt{(t_{m,1}^j)^2 + (t_{m,2}^j)^2} \right)}{2\sqrt{(t_{m,1}^j)^2 + (t_{m,2}^j)^2}}, \quad m \in \mathfrak{M},$$

where

$$t_{m,1}^j = \frac{a_{1,1}^m (u_m^j - u_{m-M}^j) + a_{1,2}^m (u_m^j - u_{m+1}^j)}{1/M},$$

and

$$t_{m,2}^j = \frac{a_{2,1}^m (u_m^j - u_{m-M}^j) + a_{2,2}^m (u_m^j - u_{m+1}^j)}{1/M}.$$

Step 4: Updating u . Update u^j by solving the linear system

$$\left(T^T T + \lambda_{T1} (R_1^T (I - \Theta) R_1 + R_2^T (I - \Theta) R_2) + \lambda_{TV} (R_1^T \Theta B^j R_1 + R_2^T \Theta B^j R_2) \right) u^j = T^T v$$

where B^j is the M^2 -by- M^2 diagonal matrix with elements b_m^j for $m \in \mathfrak{M}$ and 0 otherwise. It is worth noticing that this linear system is well posed ([8]) since the matrices appearing as a consequence of the penalization terms have strictly positive eigenvalues, and hence solving it entails no difficulties.

Step 5 - Convergence: if a previously defined convergence criterion is satisfied, the algorithm ends and our inpainted image \hat{u} is defined as u^j (in our case, this criterion was defined upon the norm $\|u^j - u^{j-1}\|$). Otherwise, the algorithm repeats from step 2.

In the next section we present some examples of the performance of the previously described inpainting method.

4 INPAINTING APPLICATION RESULTS

We begin by comparing the performance of our new approach with the isotropic order-one Tikhonov-Phillips and Curvature-Driven Diffusion methods on the previously used test image, occluded over both smooth and piecewise constant regions. A 1% Gaussian white noise was added to the grayscale image. Figure 5 depicts the occluded noisy image together with the results obtained with three different inpainting methods. As it can be clearly seen, the mixed anisotropic T1-TV inpainting outperforms the isotropic T1 method in terms of edge preservation, while it also works better than CDD inpainting on the smooth region of the image. Although this is a somewhat subjective analysis, these conclusions are supported by the corresponding *peak signal-to-noise ratios (PSNR)* shown in Table 1. The *PSNR* is defined as follows:




Figure 5: (a) Ocluded image; (b) Isotropic T1 inpainting; (c) CDD inpainting; (d) Mixed anisotropic T1-TV inpainting.

$$PSNR(\hat{u}) = 10 \log_{10} \left(\frac{1}{MSE(\hat{u}, u_0)} \right),$$

where \hat{u} , $u_0 \in \mathbb{R}^{M^2}$ are the vectors associated to the inpainted image and to the original image (unknown in real problems), respectively, and

$$MSE(\hat{u}, u_0) = \frac{1}{M^2} \|\hat{u} - u_0\|_2^2 = \frac{1}{M^2} \sum_{m=1}^{M^2} (\hat{u}_m - u_{0,m})^2.$$

	Isotropic T1	CDD	Anisotropic T1-TV
<i>PSNR</i>	20.140	35.497	36.329

Table 1: *PSNR* values for the test image (Figure 5).

Up next, we show the performance of our new method on a more realistic example. A 300×300 pixel grayscale image was occluded and contaminated with %2 Gaussian white noise. Figure 6 shows the resulting occluded noisy image, along with the

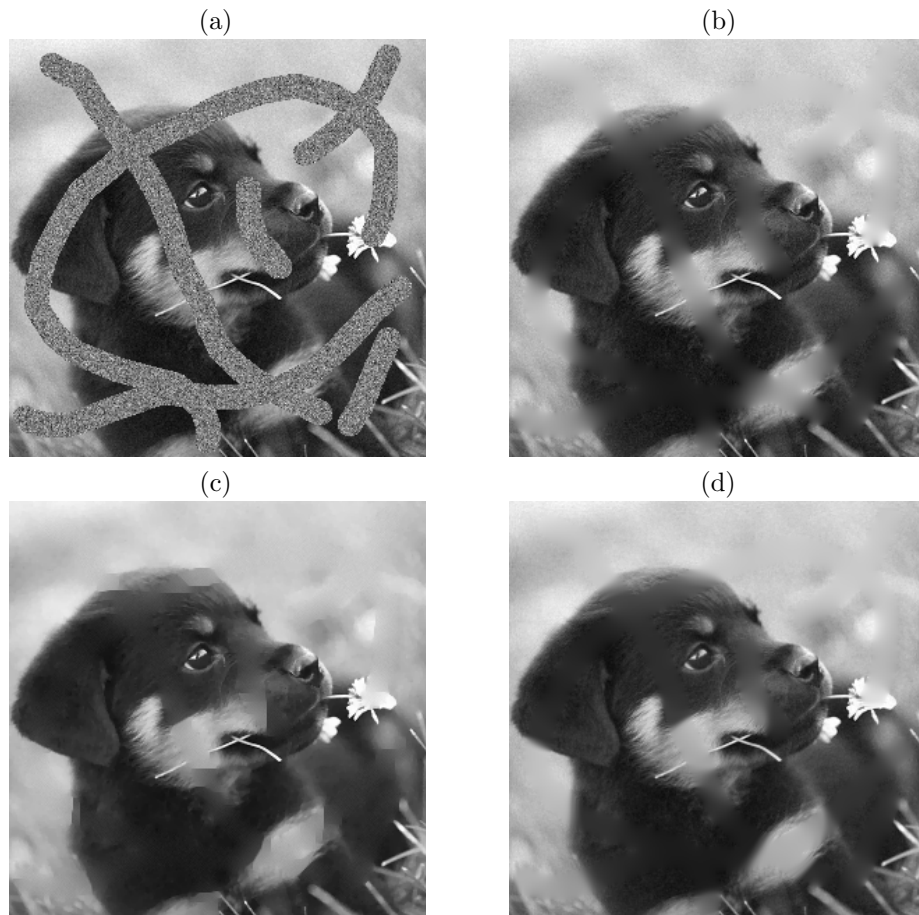


Figure 6: (a) Occluded noisy image; (b) Isotropic T1 inpainting; (c) CDD inpainting; (d) Anisotropic T1-TV inpainting.

isotropic T1, CDD and anisotropic T1-TV inpainted images. Once again we observe an improvement on the quality of the inpainting with the two-step method, which is reflected on the *PSNR* values shown on Table 2.

	Isotropic T1	CDD	Anisotropic T1-TV
<i>PSNR</i>	29.127	29.952	30.868

Table 2: *PSNR* values for grayscale image (Figure 6).

Next we show an example of the performance of our method on a color image. In this case, the inpainting process was performed separately over the red, green and blue layers of a 300×300 pixel image. Here again, a %2 Gaussian white noise was added to the occluded image. Figure 7 shows the occluded noisy image, along




Figure 7: (a) Occluded noisy image; (b) Isotropic T1 inpainting; (c) CDD inpainting; (d) Anisotropic T1-TV inpainting.

with the CDD, isotropic T1 and the anisotropic T1-TV inpainted images. A close look at Figure 7(c) shows that CDD inpainting tends to produce artificial edges, which are not produced by the combined T1-TV method. Note also that although no artificial borders are generated with the latter method, significantly better edge preservation properties are achieved with respect to the pure T1 method. Table 3 shows the obtained *PSNR* values for the three inpainting methods.

	Isotropic T1	CDD	Anisotropic T1-TV
<i>PSNR</i>	20.568	21.360	22.112

Table 3: *PSNR* values for color image (Figure 7).

Now we compare the performance of our two-step method with that of the CKS inpainting model ([3]), which is based upon Euler's elastica and curvature. This model allows for smooth inpainting, and so it provides a way to avoid the artifacts which appear using CDD. Nonetheless, this smoothing might entail losing

sharpness of inpainted edges and borders. To illustrate this, inpainting results for the same images used before are displayed on Figure 8, obtained with a 200×200 image, previously contaminated with a 2% Gaussian additive noise. It can be observed that our two-step method has better performance, as the spatial adaptivity provides a better compromise between edge preservation and smooth inpainting. Table 4 shows the *PSNR* values for the displayed restorations.

	Gray CKS	Gray A T1-TV	Color CKS	Color A T1-TV
<i>PSNR</i>	29.770	32.421	21.464	24.284

Table 4: *PSNR* values for color image (Figure 7).




Figure 8: Left: CKS inpaintings; right: Anisotropic T1-TV inpaintings.

Finally, we present an example of object removal via masking-inpainting. If one wishes to remove an object from a certain image in an “unnoticeable” way, an artificial occlusion can be made over the object, and then proceed with the inpainting. Figure 9 shows an example of the results of such a procedure over a 300×300 color image. The compromise between edge preservation and smooth inpainting of the mixed anisotropic T1-TV model can be once again appreciated here. Note that computing the *PSNR* values does not make sense in this case.

Computational cost: Although certainly there is room for improvement in the implemented algorithms, after several trials we can assure that the computational cost for the mixed inpainting approach does not increase severely with respect to the pure CDD method. Our programs were run on a PC with an Intel Core i7-3930K processor, CPU@3.20 GHz x12. For instance, for the 300×300 grayscale image

displayed on Figure 6, the CPU-time for the CDD process was 2482.4 seconds, while the CPU-time for the T1-TV model (having previously computed the pilot image u_p and the parameters) was 56.8 seconds. The CPU-time for the full algorithm, including building the pilot image, parameter estimations and auxiliar computations was 3572.4 seconds. We mention however that optimizing the algorithm was not a main goal of our work. Experimentation suggests, for instance, that a variable time-step on the diffusion process and refining parameter estimation on the T1-TV inpainting might lead to a significantly faster performance.

5 CONCLUSIONS

In this article we developed a two-step method for image inpainting. The first step consists of implementing a Curvature-Driven Diffusion process, which serves to obtain a good approximation of the gradient of the image inside the occlusion. Such an approximation is then used to construct an appropriate weighting function and an anisotropy-inducing matrix field. The second step consists of using these functions to build a mixed Tikhonov-Total Variation spatially varying anisotropic functional, whose global minimizer defines the final inpainting.

The use of a CDD inpainting process in the construction of the weighting function and the anisotropy matrix field results in the fact that the well known object-connectivity property of this method is incorporated into the mixed T1-TV model. This, along with the spatial adaptivity and anisotropy features of the mixed T1-TV method results in an inpainting model having both good object-connectivity and edge preservation properties as well as high-quality inpainting performance over smooth regions.

The performance of the method was shown through several examples, in which it produced better results than the CDD, T1 and CKS inpainting models. Furthermore, a thorough analysis of the color examples shows that while the CDD inpainting seems to sometimes produce little artifacts near edges inside the occlusions, such artifacts do not appear when using our two-step method.

Finally, it is appropriate to mention that there is room for further improvements. For instance, other means of constructing the weighting function θ and the matrix field A can be explored. Also, if there is adequate available *a-priori* information about the "expected" image, that information could be embebed into the model through the functions θ and A .

ACKNOWLEDGEMENTS

This work was supported in part by Consejo Nacional de Investigaciones Científicas y Técnicas, CONICET, through PIP 2014-2016 Nro. 11220130100216-CO and by the Air Force Office of Scientific Research, AFOSR/SOARD, through Grant FA9550-14-1-0130.

References

- [1] R. Acar and C. R. Vogel, Analysis of bounded variation penalty methods for ill-posed problems, *Inverse Problems*, **10** (1994), 1217–1229.

- [2] D. Calvetti, F. Sgallari and E. Somersalo, Image inpainting with structural bootstrap priors, *Image and Vision Computing*, **24** (2006), 782–793.
- [3] T. Chan, S. Kang and J. Shen, Euler’s elastica and curvature based inpaintings, *SIAM J. on Applied Mathematics*, **63(2)** (2002), 564–592.
- [4] T. Chan and J. Shen, Mathematical models for local nontexture inpaintings, *SIAM J. on Applied Mathematics*, **62** (2002), 259–268.
- [5] F. Ibarrola and R. Spies, Image restoration with a half-quadratic approach to mixed weighted smooth and anisotropic bounded variation regularization, *SOP Transactions on Applied Mathematics*, **1** (2014), 57–95.
- [6] J. Idier, *Bayesian Approach to Inverse Problems*, John Wiley & Sons, 2008.
- [7] F. Li, Z. Li and L. Pi, Variable exponent functionals in image restoration, *Applied Mathematics and Computation*, **216** (2010), 870–882.
- [8] G. L. Mazziari, R. D. Spies and K. G. Temperini, Mixed spatially varying L^2 -BV regularization of inverse ill-posed problems, *Journal of Inverse and Ill-posed Problems*, **23** (2014), 571–585.
- [9] R. Rockafellar, *Convex Analysis*, Princeton University Press, 1970.
- [10] L. I. Rudin, S. Osher and F. E., Nonlinear total variation based noise removal algorithms (proceedings of the 11th annual international conference of the center for nonlinear studies), *Physica D*, **60** (1992), 259–268.




Figure 9: (a) Original image; (b) Occluded (masked) image; (c) T1 inpainting; (d) CDD inpainting; (e) T1-TV inpainting.