Asociación Argentina
de Mecánica Computacional

Mecánica Computacional Vol XXXIV, págs. 3335-3349 (artículo completo)
Sebastián Giusti, Martín Pucheta y Mario Storti (Eds.)
Córdoba, 8-11 Noviembre 2016

MIXED CURVATURE-BASED DIFFUSION MODELS FOR LOCAL IMAGE INPAINTING

Francisco J. Ibarrola[a] **and Ruben D. Spies**[b]

[a]*Instituto de Investigación en Señales, Sistemas e Inteligencia Computacional, sinc(i), CONICET-UNL, Facultad de Ingeniería y Ciencias Hídricas, Universidad Nacional del Litoral, Ciudad Universitaria, CC 217, Ruta Nac. No 168, km 472.4, (3000) Santa Fe, Argentina.*
*(*fibarrola@sinc.unl.edu.ar*).*

[b]*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*).*

Keywords: Inpainting, Inverse Problems, Ill-posedness, Regularization.

Abstract. The image inpainting problem consists of restoring an image from a (possibly noisy) observation, in which data from one or more regions is 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 modeling the solution as the result of a diffusion process using the heat equation. For non-smooth images, however, such an approach is far from being satisfactory. 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 Diffusions (CDD), 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 solutions. Based upon this, we propose using CDD to introduce a-priori information into an anisotropic diffusion model that allows for both edge preservation and object connectivity while precluding the staircasing effect that TV-based methods entail. Comparisons between the results of the implemented models will be illustrated by several computed examples, along with performance measures.

1 INTRODUCTION

The image inpainting problem consists of restoring an image from an occluded and possibly noisy observation of it, i.e. data from one or more regions are missing. There are a few proposed models to perform this task and although some of them perform reasonably well in certain types of images, all of them are far from being totally satisfactory. If the image is expected to be smooth, for instance, inpainting can be reasonably performed by means of a Bayesian approach and a maximum a posteriori computation Calvetti et al. (2006), while for non-smooth images such an approach is far from being satisfactory. Although the introduction of anisotropy ideas to the latter methodology is known to produce better results for slim occlusions Calvetti et al. (2006), the quality of the restoration decays as the occluded regions widen. also Total Variation (TV)and inpainting models based on PDE diffusion equations can be used whenever edge restoration is a priority. Recently variants of these methods, such as Curvature-Driven Diffusion (CDD) inpainting Chan and Shen (2002) which takes into account curvature into the diffusion coefficient, have resulted in inpainted images with very good edge preservation and object connectivity properties. The CDD approach, nonetheless, does not produce satisfactory results in regions where the image is smooth, as it tends to produce piecewise constant restorations, which is a reminiscent manifestation of its TV origins.

In this article we present a two-step inpainting process consisting of a first in painting round for building a pilot image from which to infer *a-priori* structural information on the image's gradient via CDD and a second step where the final inpainting is performed via mixed L^2-anisotropic TV regularization. We present a few results showing the improvement of our two-step approach over all preexisting inpainting methods.

Our (in principle grayscale) image is defined by a function $u : \Omega \subset \mathbb{R}^2 \to [0, 1]$, where $u(x, y)$ represents the light intensity of the point (x, y) ($u = 1$ being white and $u = 0$ being black). The occlusion will be denoted by D while $v \doteq u|_{\Omega\setminus D}$ shall denote the known part of u.

Before describing our two-step method, we briefly recall three of the most traditional inpainting methods.

1.1 Tikhonov-Phillips Inpainting

The Tikhonov-Phillips regularization method of order 1 (T1) can be used to performed a basic inpainting. With this method, the inpainted solution is defined as

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

where $\mathcal{T} : L^2(\Omega) \to L^2(\Omega \setminus D)$ is the D-occlusion operator, $\mathcal{T}u = u\ |_{\Omega\setminus D}$ and $\lambda > 0$ is an appropriately chosen regularization parameter. 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)). This lack of edge preservation can then be improved by appropriately modifying the penalizing term in (1), for instance by means of the introduction of introduce an anisotropy matrix field A Calvetti et al. (2006), which attenuates penalization on on directions of large gradients. In this case, the inpainted image is defined as

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

The construction of A requieres prior information of the gradient inside the occlusion. Ways of constructing A will be presented later. 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 Calvetti et al. (2006). Although this method significantly improves the performance os its isotropic counterpart, it becomes inappropriate as occlussions widen (see Figure 1(c)).

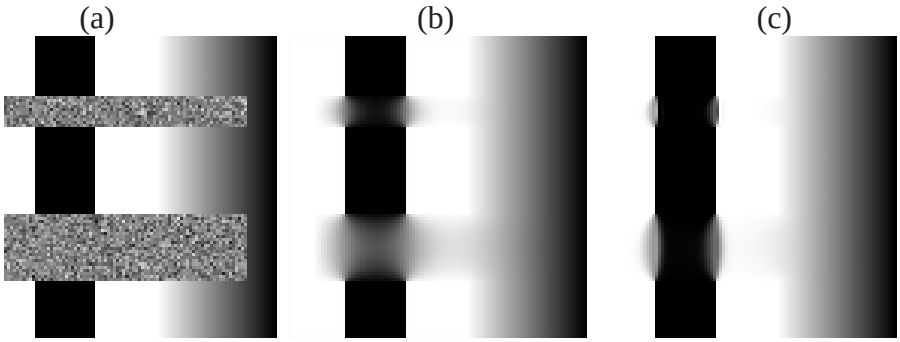

Figure 1: (a) Occluded image; (b) T1 inpainting; (c) Anistoropic T1 inpainting.

1.2 Total Variation inpainting

Total variation Acar and Vogel (1994), Rudin et al. (1992) inpainting is defined by:

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

The Euler-Lagrange equations show that the solution of (3) is also the steady-state solution of the following diffusion PDE (see Li et al. (2010))

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

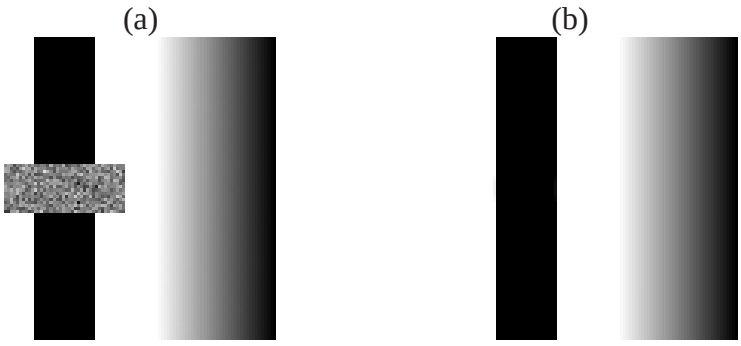

Figure 2: Occluded image (a); TV inpainting (b).

The property of being able to link similar objects on opposite sides of an occlusion is referred to as *object connectivity*. Although TV inpainting improves object connectivity it has two main draw backs. One one side, it tends to produce piecewise constant solutions (see Figure 2) and on the other hand the results strongly depend on the width of the occlusion (see Figures 2(a) and 2(b)).

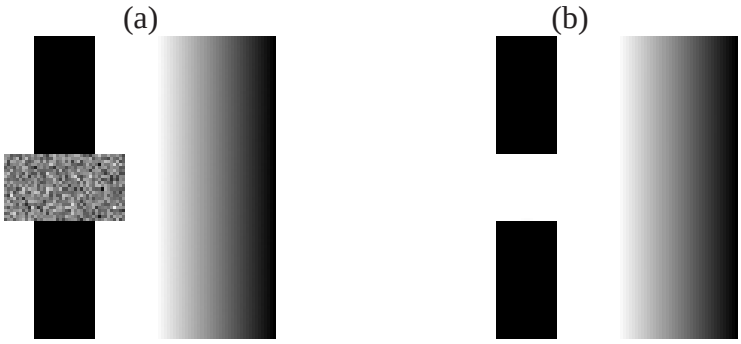

Figure 3: (a) Occluded image; (b) TV inpainting.

Although this puts into evidence that any inpainting process can entail a high level of subjectivity, 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

Roughly speaking, an *isophote* of an image u can be thought of as a 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 occlusions since this will result in the "reconnection" of edges set apart by the occlusion. In fact, in Figure 3(b) we see that the original isophotes (the edges of the black bar) do not result connected as we (or at least most of us!) 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 Chan and Shen (2002) to include the curvature into the diffusion model (4). By letting $\hat{D} = \frac{1}{|\nabla u|}$, equation (4) restricted to the occlusion reads

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

In Chan and Shen (2002) the diffusion coefficient $\hat{D}$ was redefining as $\hat{D} = \frac{g(|\kappa|)}{|\nabla u|}$, where $g : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ 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

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

Since the curvature of the level lines is $\pm\infty$ at the corners of the TV inpainted image in Figure 3(b), is is clear now that 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, as mentioned before, being this method a byproduct of TV regularization, it tends to produce piecewise constant restorations, and therefore it does not produce completely satisfactory results over smooth regions (see Figure 4).

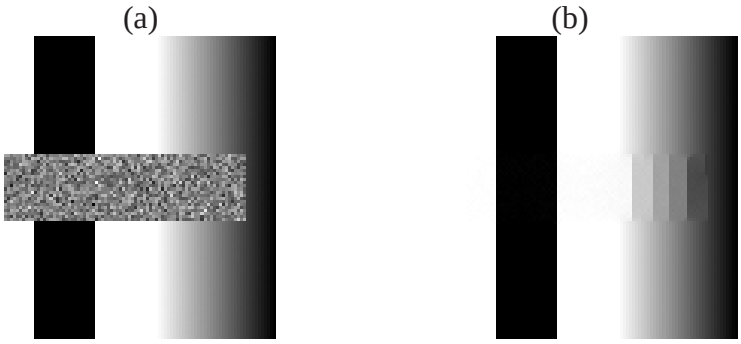

Figure 4: (a) Occluded image; (b) CDD inpainting.

The CDD inpainting model (5) can also take into account the presence of noise outside the occlusion, for instance, by using TV regularization outside the occluded region. The resulting diffusion PDE takes then the form

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

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 which simultaneously complies with edge preservation, object connectivity and good inpainting performance over smooth regions.

2 COMBINED, TWO STEP CDD + T1-TV INPAINTING

Regularization methods can be combined by using spatially-varying weighted averages of two or more penalizers Mazzieri et al. (2014) to extract the the good characteristic properties of each one. We shall consider here 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. In this case the inpainted image is then be defined as the minimizer of

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

where $\lambda_{T1} > 0$ and $\lambda_{TV} > 0$ are appropriately chosen 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 $\theta = 0$ leads to pure anisotropic T1 regularization, while $\theta = 1$ leads to pure anisotropic TV regularization.

Thereare several ways of constructing the matrix field $A : \Omega \to \mathbb{R}^{2\times 2}$ used for introducing anisotropy. One way can be found in Calvetti et al. (2006), where it is built by using the gradient field of an *a-priori* estimation $u_p(x, y)$ of $u(x, y)$, as follows. First a continuous and decreasing function $h : \mathbb{R}^+_0 \to (0, 1]$, satisfying $h(0) = 1$ and $\lim_{t\to\infty} h(t) = 0$ is defined, to determine the eigenvalues of $A(x, y)$, which is constructed as follows:

$$A(x, y) = I - (1 - h(|\nabla u_p(x, y)|)) \left[\frac{\nabla u_p(x, y)}{|\nabla u_p(x, y)|} \right] \left[\frac{\nabla u_p(x, y)}{|\nabla u_p(x, y)|} \right]^{\mathrm{T}}. \tag{8}$$

As a consequence, A has the following important 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

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

With the required properties on h, the eigenvalues $\sigma_1(x,y)$ are 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. To overcome the disadvantage of lack of edge preservation presented for most methods, we shall use CDD to construct the gradient field estimation ∇u_p. Hence we shall build both the weighting function θ and the matrix field A, based upon the gradients of an *a-priori* CDD inpainting of the image, which as previously noted, also favors object connectivity. In particular we take $\theta(x,y) = w(|\nabla u_p(x,y)|)$, where w is an appropriately chosen increasing function, with $0 \leq w(t) \leq 1$ $\forall t \in \mathbb{R}_0^+$. It is important to remark that 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 new full inpainting process can then be stated as follows:

Step 1: CDD inpainting. Perform a CDD inpainting to obtain a first pilot image $\tilde{u}_p$ 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, for instance, a low-variance Gaussian kernel.

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

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

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 A and θ as above to build the mixed weighted T1-TV anisotropic functional and compute the restored image as

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

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 $u_{M(l-1)+m} = U_{m,l}\ \forall l, m = 1, 2, \ldots, M$. For better understanding we will often identify $u_{M(l-1)+m}$ with $u(x, y)$.

3.1 The CDD method

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

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

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

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

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

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

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 Chan and Shen (2002)), 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

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

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

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

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

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

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 Ibarrola and Spies (2014)).

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 $\phi : \mathbb{R} \to \mathbb{R}$ 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 (Rockafellar (1970)) that there exists a function ψ satisfying the following duality relation

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

and therefore

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

where

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

and

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

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 $\Theta \doteq diag(\theta_m)_{M^2\times M^2}$ and $S \doteq diag(s_m)_{M^2\times M^2}$. 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

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

It can be shown (Ibarrola and Spies (2014)) that

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

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

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

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

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

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

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

Although details on de derivation of equation (21) can be found in Idier (2008), we give here a brief sketch of the proof. Let $f : \mathbb{R}^2 \to \mathbb{R}$ 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

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

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 (Mazzieri et al. (2014)) 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:

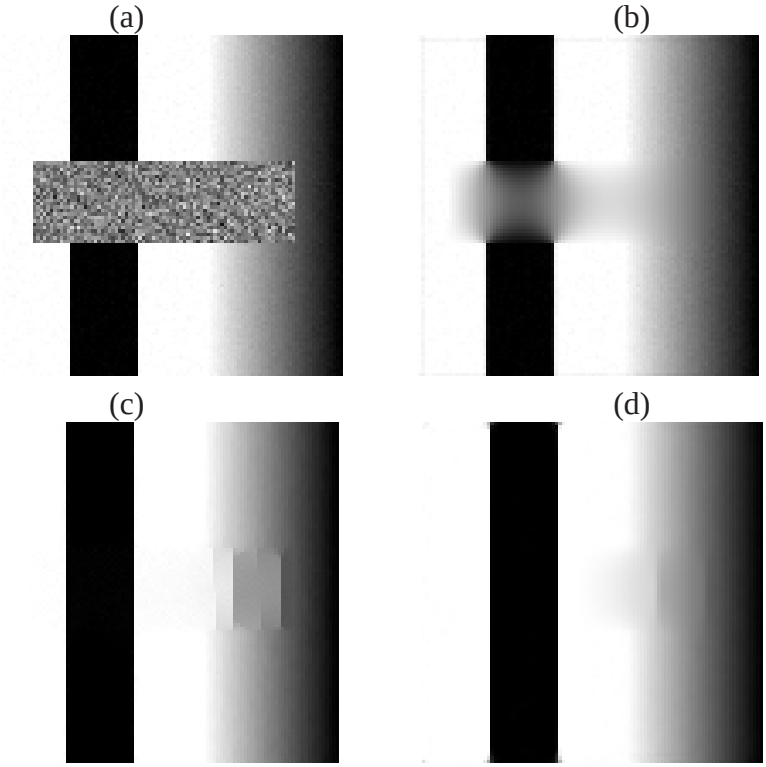

Figure 5: (a) Occluded image; (b) Isotropic T1 inpainting; (c) CDD inpainting; (d) Mixed anisotropic T1-TV inpainting.

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

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
PSNR	20.140	35.497	36.329

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

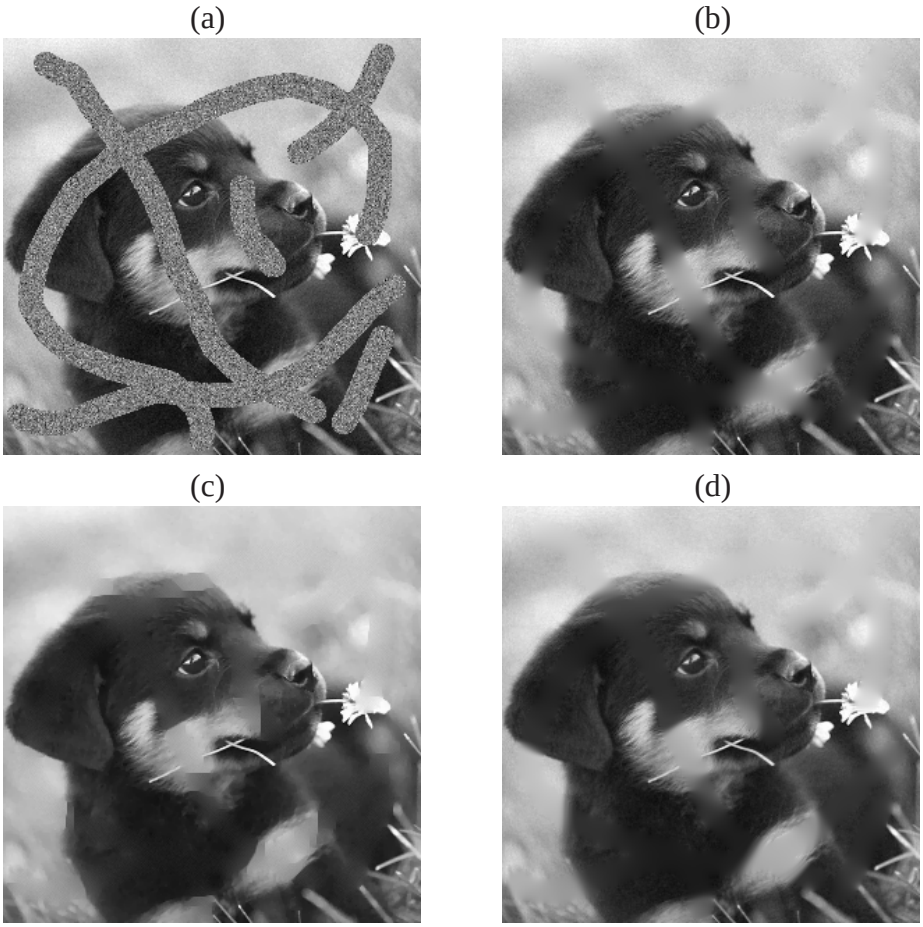

Figure 6: (a) Occluded noisy image; (b) Isotropic T1 inpainting; (c) CDD inpainting; (d) Anisotropic T1-TV inpainting.

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 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
PSNR	29.127	29.952	30.868

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

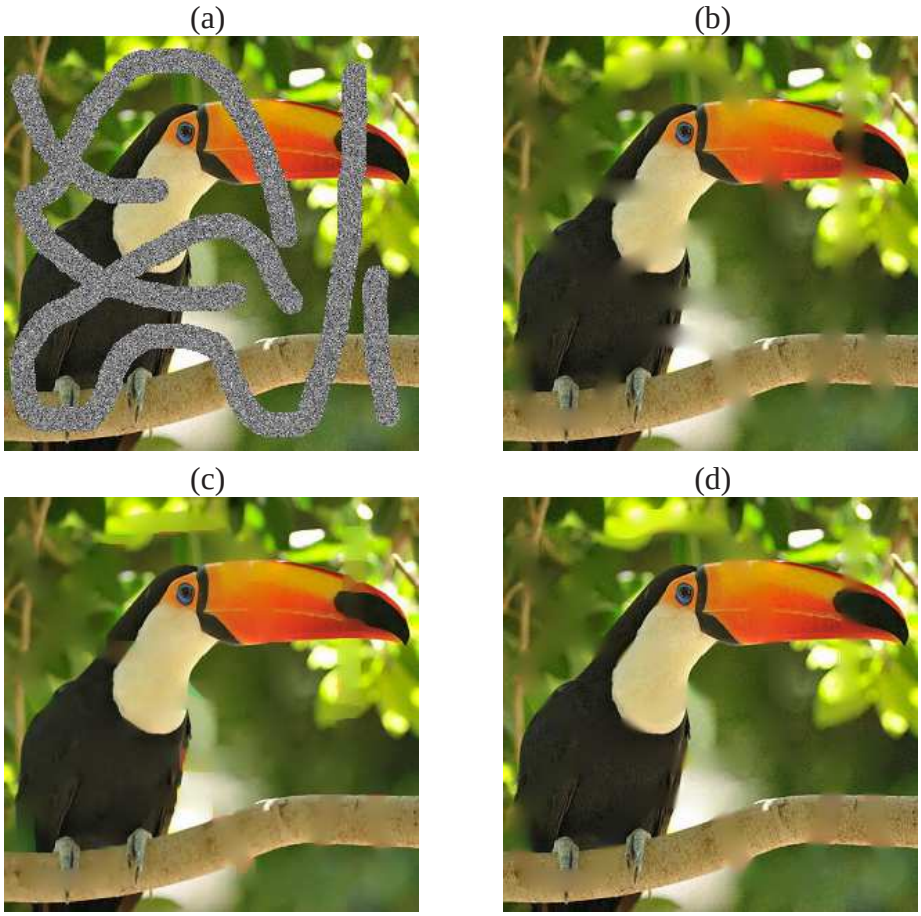

Figure 7: (a) Occluded noisy image; (b) Isotropic T1 inpainting; (c) CDD inpainting; (d) Anisotropic T1-TV inpainting.

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 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
PSNR	20.568	21.360	22.112

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

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 and T1 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

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

Calvetti D., Sgallari F., and Somersalo E. Image inpainting with structural boostrap priors. *Image and Vision Computing*, 24:782–793, 2006.

Chan T., Kang S., and Shen J. Euler's elastica and curvature based inpaintings. *SIAM J. on Applied Mathematics*, 63(2):564–592, 2002.

Chan T. and Shen J. Mathematical models for local nontexture inpaintings. *SIAM J. on Applied Mathematics*, 62:259–268, 2002.

Ibarrola F. and Spies R. Image restoration with a half-quadratic approach to mixed weighted smooth and anisotropic bounded variation regularization. *SOP Transactions on Applied Mathematics*, 1:57–95, 2014.

Idier J. *Bayesian Approach to Inverse Problems*. John Wiley & Sons, 2008.

Li F., Li Z., and Pi L. Variable exponent functionals in image restoration. *Applied Mathematics and Computation*, 216:870–882, 2010.

Mazzieri G.L., Spies R.D., and Temperini K.G. Mixed spatially varying $l2$-bv regularization of inverse ill-posed problems. *Journal of Inverse and Ill-posed Problems*, 22, 2014.

Rockafellar R. *Convex Analysis*. Princeton University Press, 1970.

Rudin L.I., Osher S., and E. F. Nonlinear total variation based noise removal algorithms (proceedings of the 11th annual international conference of the center for nonlinear studies). *Physica D*, 60:259–268, 1992.