ADAPTIVE REST CONDITION POTENTIALS: FIRST AND SECOND ORDER EDGE-PRESERVING REGULARIZATION

Mariano Rivera & José Luis Marroquín

Comunicación Técnica No I-02-22/21-10-2002
(CC/CIMAT)

Adaptive Rest Condition Potentials: First and Second Order Edge-Preserving Regularization

Mariano Rivera and Jose L. Marroquin
Centro de Investigacion en Matematicas A.C.
Apdo. Postal 402, Guanajuato, Gto. Mexico 36020
{mrivera,jlm}@cimat.mx
http://www.cimat.mx/∼mrivera

September 23, 2002

Abstract

A new regularization formulation for inverse problems in computer vision and image processing is introduced, which allows one to reconstruct second order piecewise smooth images, that is, images consisting of an assembly of regions with almost constant value, almost constant slope or almost constant curvature. This formulation is based on the idea of using potential functions that correspond to springs or thin plates with an adaptive rest condition. Efficient algorithms for computing the solution, and examples illustrating the performance of this scheme, compared with other known regularization schemes are presented as well. *Keywords. Edge-preserving regularization, image restoration, segmentation, anisotropic diffusion.*

1 Introduction

Since the seminal works in Refs. [1]-[4] about image (or surface) restoration with discontinuities, in recent years, several methods for Edge-Preserving Regularization (EPR) for inverse problems in computer vision and image processing, have been published [5]–[27]. These methods have demonstrated their performance in detecting outliers in the data and reconstructing piecewise smooth images. The definition of piecewise smooth, however, has in most cases meant "almost piecewise constant", which means that the image can be represented as an assembly of regions such that inside them the gradient is close to zero. In the regularization framework, given the observed image g, the regularized solution f is computed as the minimizer of an energy functional U. Although convex potentials guarantee convergence to its global minimum [6]–[12], sharper edges are reconstructed with non-convex potentials that grow at a slower rate than quadratic ones. Provided a good initial guess for f is given, efficient algorithms for computing a local minimum have been reported in the literature [13]–[22]. In spite of the success of robust regularization methods, there are still important open problems; in particular, the definition of piecewise smooth images has not been extended successfully to include regions with almost constant slope (second order smoothness). As a result, regions with constant slope are reconstructed with a "staircase" effect.

The purpose of this paper is to introduce a new formulation for energy potentials that allows one to reconstruct images with second order piecewise smoothness. In addition, efficient algorithms for computing the solution are presented. The organization of the paper is as follows: section 2 presents a review of the EPR techniques based on robust potentials. In order to clarify the behavior of first order robust regularization an analogy with a Weak Spring System (WSS) is used. We show that this model has limitations for representing potentials for high order EPR.

The third section introduces a new formulation for potentials based on the paradigm of the adaptive rest condition (ARC). For first order potentials, the corresponding analogous model is a spring system with

adaptive rest condition (SARC); we show that for this case, there is an equivalence between WSS and SARC potentials, although the SARC formulation allows an increased flexibility. In section 4, we extend the ARC formulation to second order potentials, which do not have an equivalent representation in terms of known models. We call this model the Plate System with Adaptive Rest Condition (PARC), and show that it allows one to correctly reconstruct piecewise planar images. Also in this section we present efficient minimization algorithms for the corresponding energy functions. In section 5 we present experimental results for the new systems, both with synthetic and real images, and compare their performance with other EPR techniques. Finally, our conclusions are given in section 6.

2 Robust Regularization

2.1 Statement of the Problem

The problem of reconstructing an image f from noisy and degraded observations g given the following model of the observations:

$$g = F(f) + \eta, \tag{1}$$

where η is additive noise and F is (in general) a non-linear operator that is assumed to be known, is an *ill posed* problem. Therefore, regularization of the problem is necessary. This means that, *prior* information or assumptions about the structure of f need to be introduced in the reconstruction process. The regularized solution f^* is computed by minimizing an energy functional U:

$$f^* = \arg\min_f U(f)$$

where U is of the form:

$$U(f) = D(f, g) + \lambda R(f), \tag{2}$$

The first term in (2) establishes that the reconstructed f should be consistent with the data g and the second term imposes a penalty for violating the prior assumptions about f, e.g., piecewise smoothness. The relative contribution of each term to the global energy is controlled by the positive parameter λ.

2.2 The Homogeneous Spring System

In the framework of Bayesian regularization, the data term in (2) is chosen as the negative log-likelihood and the prior constraints are incorporated in the form of a prior MRF model for f [1][3][28], so that the regularization term R in (2) takes the form of a sum, over the cliques of a given neighborhood system, of a set of "potential functions" supported on those cliques. One may take for instance as the neighborhood N of a pixel r its 8 closest neighbors:

$$N_r = \{s : |r - s| < 2\}$$

and cliques of size 2 $\langle r, s\rangle$ that correspond to horizontal, vertical and diagonal pixel pairs, where $r = (x, y)$ represents a site in the pixel lattice L. A quadratic regularization energy is obtained by assuming that η corresponds to Gaussian noise and choosing quadratic potentials over the first neighbor pairs:

$$U_H(f) = \sum_r \left\{ |F(f)_r - g_r|^2 + \frac{\lambda}{2} \sum_{s \in N_r} d_{rs} \, |f_r - f_s|^2 \right\} \tag{3}$$

where the constant d_{rs} is equal to the inverse of the distance between the pixels r and s ($d_{rs} = 1/|r - s|$). Functional (3) corresponds to the internal energy of the physical model of a Homogeneous Spring System. This model is equivalent to a system of particles located at the sites of the pixel lattice, so that the vertical position of each particle is represented by the gray level of the corresponding pixel. Eq. (3) corresponds to the energy of the complete system where (when F is the identity) each particle f_r is connected by means of

springs with the observation g_r and with its neighboring particles. The cost functional (3) does not preserve edges and will produce an over-smoothing of the real edges of the image.

2.3 The Weak Spring System: Robust Regularization

To alleviate that problem, there have been proposed potential functions for the regularization term that allow edge preservation, based on the idea of a breakable spring, that is, if the potential energy of a spring exceeds a given threshold θ, then the spring must be broken [1][3] or weakened [5][13][16][17]. To achieve this behavior, an auxiliary variable ω than acts as edge (outlier) detector is introduced; then the potential takes the form:

$$\rho(f_r - f_s, \omega_{rs}) = (f_r - f_s)^2 \omega_{rs} + \Psi(\omega_{rs}), \tag{4}$$

where ω_{rs} is associated to each pixel pair (r, s), and Ψ is a potential function that controls an over-detection of edges. In the case of the breakable spring model [1], ω_{rs} only takes the values $\{0, 1\}$; on other hand, in the case of the WSS model, $\omega_{rs} \in [0, 1]$, and is set close to 1 for $(f_r - f_s)^2 < \theta$ (where θ is a given threshold) and less that one otherwise. Black and Rangarajan [13] have shown that the potentials of the weak spring model correspond to the cost function for robust M-estimators. These potentials are, in general, non-convex and grow at a slower rate than the quadratic ones. This method is capable of finding the significant missing data of a noisy image and performing an edge-preserving restoration. Furthermore, the explicit outlier detection formulation allows one to incorporate additional constraints about the structure of the edge reject field ω [13][25][26]. For instance, one can penalize the "thickness" and the discontinuities on the edges, at the expense of an additional computational cost.

2.4 The Weak Thin Plate Model

The thin plate model [2][5][28] is obtained when one uses as potentials, squares of finite difference approximations to second derivatives:

$$\Delta^2 f_r = f_q - 2f_r + f_s. \tag{5}$$

The computation of $\Delta^2 f_r$ involves cliques of size 3 $\langle q, r, s \rangle$ that correspond to horizontal, vertical and diagonals pixel triads (see figure 1).

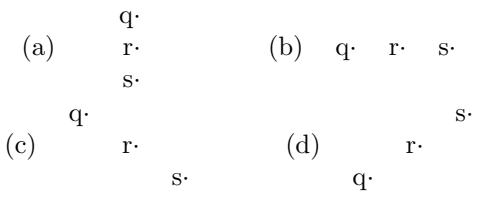

Figure 1: Cliques with triads of pixels

One could use the weak potential

$$\rho(\Delta^2 f_r, \omega_{qrs}) = \left[\Delta^2 f_r\right]^2 \omega_{qrs} + \mu\Psi\left(\omega_{qrs}\right), \tag{6}$$

as a second order edge-preserving potential; however, the results are not completely satisfactory, even for the reconstruction of piecewise constant images, where the first order model presents an excellent performance. The observed effect consists in the "ramping" or interpolation of first order discontinuities. In order to compute a second order solution, Geman and Reynolds [5] proposed to use the reconstruction computed with the first order model [using $\rho(f_r - f_s, \omega_{rs})$] as the starting point for the second order model. This improves the results, but still presents some problems, because the outliers for the first order model (jumps) do not correspond to the outliers for the second order one (large curvatures). Thus, the weak second order model does not work well in the edges defined by jumps in the gray level (see section 5).

3 The Adaptive Rest Condition Potentials (ARC-Potentials)

3.1 Spring System with Adaptive Rest Condition

In this section we introduce the potentials for EPR that are based on the paradigm of Adaptive Rest Condition (ARC). The system we are proposing is based on the idea of using quadratic potentials with a non–zero adaptive rest condition, which is expressed as the product of a function ϕ of the gradient times an "edge variable" l. In the first order case, the potentials for springs with ARC (SARC potentials) are thus:

$$\rho_{sarc}(f_r, f_s) = |f_r - f_s - \phi(f_r - f_s)l_{rs}|^2 \tag{7}$$

and the complete energy is:

$$U_{sarc}(f, l) = \sum_r \{|F(f)_r - g_r|^2 + \frac{\lambda}{2} \sum_{s \in N_r} [\rho_{sarc}(f_r, f_s) + \Phi(l, r, s)]\} \tag{8}$$

where $\Phi(l, r, s)$ is a term that penalizes the appearance of edges in the solution.

This energy will have the correct behavior if $\phi(f_r - f_s)l_{rs}$ is close to zero for those values of f_r, f_s one wants to smooth out (i.e., in the interior of regions with constant intensity), and $\phi(f_r - f_s)l_{rs} \approx f_r - f_s$ for those values that are considered edges. We have 4 possible strategies for choosing ϕ and l:

1. Explicit Line Process (SARC-EL). This corresponds to the choice $\phi(t) = t$, thus (7) can be written as $|t_{rs}(1 - l_{rs})|^2$. If $\Phi(l, r, s)$ is a convex function of l_{rs}, this case is equivalent to the WSS model; for instance, choosing $\Phi(l, r, s) = \mu l_{rs}^2$, one gets the Ambrosio–Tortorelli potential [19]. Note that in this case, the optimal l is always in the interval $[0, 1]$.

2. Implicit Line Process (SARC-IL). This corresponds to the choice $l_{rs} = 1$ for all r, s, and $\phi(t)$ shaped as in Fig. 2. This may be obtained by setting $\phi(t) = t - \rho'(t)/2$, where $\rho(t)$ is a robust potential. Note that with this choice, SARC-IL is equivalent to the implicit WSS model [13], with potentials equal to $(\rho'/2)^2$.

3. Line Process Only (SARC-LP). This corresponds to the choice $\phi(t) = 1$. Choosing $\Phi(l, r, s) = \mu l_{rs}^2$, one gets a formulation similar to that in [14][15]. Note that l approximates the gradient of f (at least close to the detected edges); hence, if a penalization term for the gradient of l is added (as in [18] [24]), this becomes, in some sense, a second order formulation.

4. General Case (SARC-G). This is a new formulation, different from the known first order EPR schemes. It corresponds to the choice $\Phi(l, r, s) = \mu l_{rs}^2$, and $\phi(t)$ chosen as in the SARC-IL case. A robust potential that gives good results is:

   $$\rho(t) = \varepsilon t^2 + 1 - \frac{(1 - \varepsilon)}{k} \exp(-kt^2)$$

   where ε is a small constant used to penalize very large gradient values and k is a scale parameter. This gives:

   $$\phi(t) = (1 - \varepsilon)t(1 - \exp[-kt^2]) \tag{9}$$

   Note that for small values of $f_r - f_s$, l will be close to zero (by effect of the penalization term), whereas for large values it will be close to one (since $\phi(t) \approx t$). This formulation has in general better performance than the WSS, and produces cleaner edge maps; Fig. 3 shows an experiment that illustrates this fact. Note that for large values of $f_r - f_s$, $\phi(f_r - f_s)$ will approximate the gradient of f, and l will act as an edge indicator variable.

As in the case of the half-quadratic algorithms, the minimization of the regularized functional, $U_{SARC}(f, l)$ in the SARC-G case, is performed in a two–step procedure: given an initial estimation f^0 for f, repeat until convergence:

1. Step 1. Minimize $U_{SARC}(f, l)$ with respect to l keeping f fixed.
2. Step 2. Minimize $U_{SARC}(f, l)$ with respect to f keeping l fixed.

Step *1* gives a closed–form solution for the optimal l:

$$l_{r,s} = \frac{\mu\phi(f_r - f_s)}{\mu + \phi^2(f_r - f_s)} (f_r - f_s), \quad \text{for all} \langle r, s\rangle \in L.$$

The minimization of step *2* is relatively difficult, since $U_{SARC}(f, l)$ is a non–convex function of f. However, it is not necessary to carry out a complete minimization, and it is enough to guarantee that U_{SARC} decreases after step *2*.

In our implementation, this is achieved by performing a single iteration of a modified Gauss–Newton scheme. The Gauss–Newton's algorithm is given by $f^{t+1} = f^t + \delta^t$, where δ^t represents the update of f in step t and it is computed as the solution of the linear system

$$\tilde{H}(f^t)\delta^t = -G(f^t),$$

where $G(f^t)$ and $\tilde{H}(f^t)$ represents the gradient and an approximation of the Hessian $H(f^t)$ of U_{SARC} in the point (f^t, l^t), respectively. To solve the linear system for δ we use the Linear Conjugated Gradient algorithm. In practice, we perform 3 iterations of in each step. If the step δ^t does not reduce the energy, we apply the backtracking method presented in subsection 4.2.

This algorithm converges faster than the classical EPR scheme (e.g. the algorithms based on an alternated minimization reported in [5][16]). The reason for this is that in the classical scheme ($\phi(t) = t$) the border indicators $l_{r,s}$ take their extreme (convergence) values $\{0, 1\}$ only if $|f_r - f_s| = 0$ (in which case $l_{r,s} = 0$) or $|f_r - f_s| > \mu$ (in which case $l_{r,s} = 1$). When $|f_r - f_s| \in (0, \mu)$, $l_{r,s}$ takes intermediate values which slow down the convergence of the method. In our case, by contrast, the edge variable $l_{r,s}$ takes intermediate values in (0,1) only in a very small interval of values of $|f_r - f_s|$, and as a result, the system converges faster and to a cleaner solution[21].

The SARC-G (and also SARC-EL) formulation has the additional advantage of allowing the introduction of explicit constraints on the detected edges [27], e.g., using potentials that favor sharp changes in the l field in directions perpendicular to the edges; one may use, as in [25][26], potentials of the form:

$$\Phi(l, r, s) = \mu l_{rs}^2 + \rho_l(l_{qr} - l_{rs}), \tag{10}$$

where ρ_l is in general a robust potential and site q is a neighbor to sites r, s, as in Fig. 1. In the present (first order) case, the effect of introducing these potentials is not very noticeable and increases the computational complexity. In the second order case, however, it has a significant effect (see sections 4 and 5).

(a) (b)

Figure 2: (a) A typical rest condition ϕl (solid line) and the residual error $t = f_r - f_s$ (doted line). (b) ARC-Potential corresponding to $|t - \phi l|^2$, with ϕl plotted in panel(a). Doted line: quadratic potential t^2

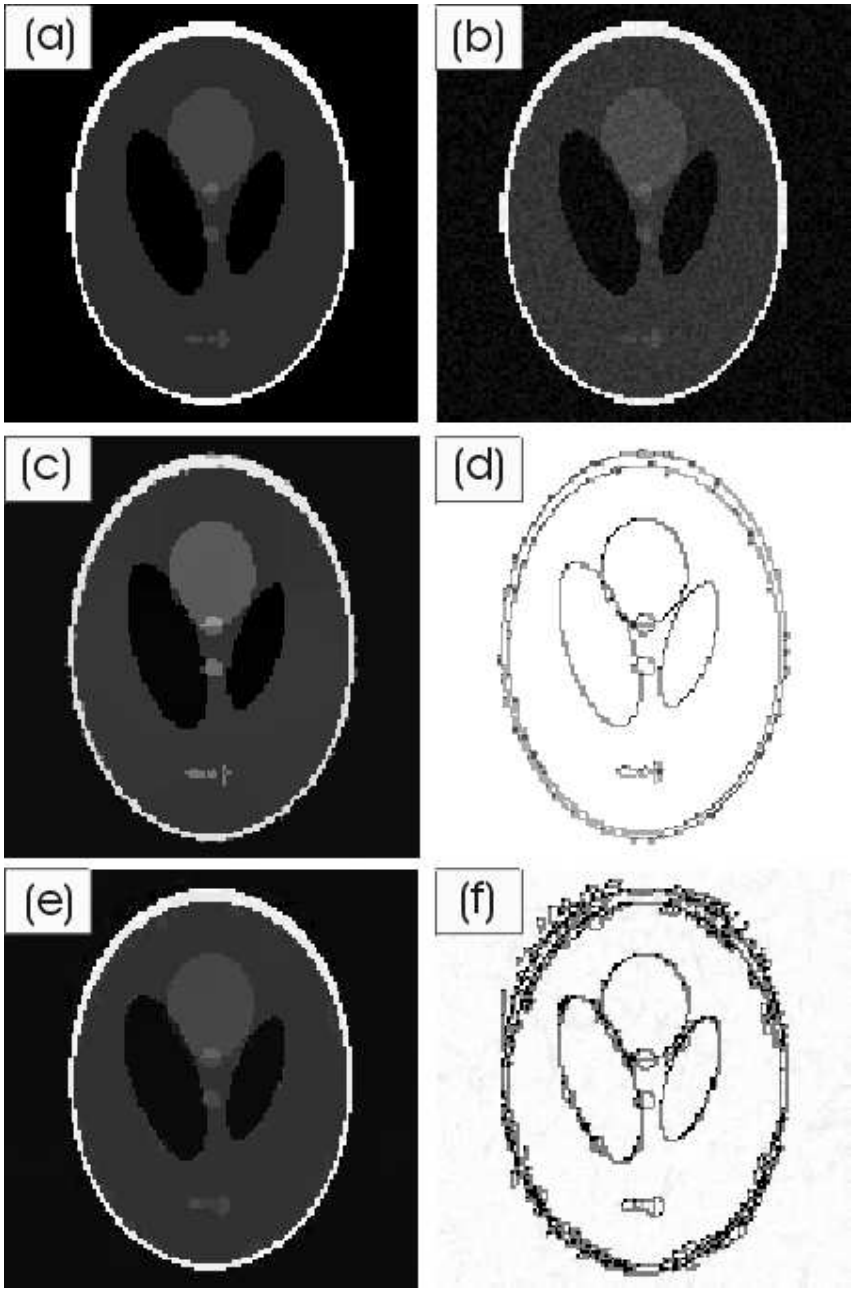

Figure 3: Filtered images with a rest condition that combines an indicator variable l and a function $\phi(t)$ given by (9). (a) Original image. (b) Noisy observations. (c) and (d) Filtered image and edge map obtained with the SARC-G model. (e) and (f) Corresponding results obtained with the WSS model (i.e. $\phi(t) = t$).

4 Thin Plate System with Adaptive Rest Condition (PARC)

The greatest advantage of the ARC formulation is that it can be easily extended to the second order case, by defining thin plate potentials with adaptive rest condition. These PARC regularization potentials have the property of not just adapting their stiffness, but also changing their behavior to SARC potentials at the edges of almost constant regions. This represents a significant advantage over the half-quadratic plate model based on robust potentials[4][5][13][14] [15][16].

The generalization of the SARC model to the second order case is based on the observation that the second order (plate) potential (5) can be written as the difference between 2 spring potentials:

$$\Delta^2 f_r = \Delta^+ f_r - \Delta^- f_r \tag{11}$$

where

$$\Delta^+ f_r = f_q - f_r \quad \text{and} \quad \Delta^- f_r = f_r - f_s \tag{12}$$

(see figure 1). Using an ARC for each one of these first order potentials, one gets the general PARC potential as:

$$\rho_{parc}(f_q, f_r, f_s) = \left|\Delta^+ f_r - \phi(\Delta^+ f_r) l_{qr} - (\Delta^- f_r - \phi(\Delta^- f_r) l_{rs})\right|^2, \tag{13}$$

where the function $\phi()$ and the edge indicator variables l are defined as in the SARC case (Eq. (9)). The complete energy is now:

$$U_{parc}(f, l) = \sum_r \{|F(f)_r - g_r|^2 + \frac{\lambda}{2} \sum_{q,s \in N_r} [\rho_{parc}(f_q, f_r, f_s) + \Phi(l, r, s) + \Phi(l, q, r)]\} \tag{14}$$

Note that this energy will have the desired behavior: in the interior of regions of constant slope, l_{qr} and l_{rs} will be close to zero, by effect of the Φ penalization terms; in the vicinity of edges between uniform regions, one of the rest conditions will be activated and the PARC potential will behave like a spring, and at edges between regions of non–zero constant slope, both ARC's will become active.

As in the SARC case, there are different strategies for choosing ϕ and l, which generate different second order EPR schemes; in this case, however, they are not equivalent to any known model. We consider 2 cases:

1. Explicit Line processes (PARC-EL), which corresponds to the choice $\phi(t) = t$ and $\Phi(l) = \mu l^2$.
2. Implicit Line processes (PARC-IL), which corresponds to the choice $l_{rs} = 1$ for all r, s, and ϕ given by (9).

We now explain the minimization algorithms that are used in each case.

4.1 Half-Quadratic Coupled Minimization Algorithms for PARC-EL Models

In this case, if F is a linear operator, U_{parc} is quadratic in f for a given l and quadratic in l for a given f, so that standard half–quadratic techniques that alternatively minimize f and l may be used. The PARC-EL model can also incorporate potentials that penalize the structure of the auxiliary variable, as in the case of the first order potentials, so that Φ is given by (10). In this case, the algorithm becomes more complicated, since now the minimization of l for a given f is itself a half–quadratic minimization problem, so that the solution of the complete system involves an iterative process in which 3 uncoupled systems of linear equations have to be solved at each iteration. In practice, however, it is not necessary to perform the complete quadratic minimization at each step; we have found very good results by alternating just one Gauss–Seidel iteration for each system at each global iteration. This scheme, with the Geman–McClure potential for ρ_l in (10) is labelled PARC-EL$^+$ in the experiments described in section 5.

4.2 Adaptive Non-Linear Conjugate Gradient Algorithm (ANLCG) for PARC-IL Models

In this case, U_{parc} is differentiable, but non-quadratic in f, so a non-linear optimization algorithm needs to be used to find a minimum. We propose here a modification to the Non-Linear Conjugate Gradient Algorithm (ANLCG), in which the step size is adaptively varied. Additionally, in order to accelerate the convergence rate, the algorithm introduces inertia in the descent. The algorithm is:

ANLCG

Set $n = 1, \beta_0 = 0, f_0$ equal to an initial guess, and $g_0 = \nabla U(f_0)$

Repeat until $|g_n| < \varepsilon$:

1. $s_n = -g_n + \beta_n s_{n-1}$
2. Compute the step α_n such that ensures energy reduction. i.e. $U(f_n + \alpha_n s_n) < U(f_n)$ (see below)
3. $f_{n+1} = f_n + \alpha_n s_n$,
   $n = n + 1$.
4. $g_n = \nabla U(f_n)$
5. $\beta_n = \max\left\{0, \frac{g_n^T (g_n - g_{n-1})}{g_{n-1}^T g_{n-1}}\right\}$,

where $\varepsilon \in (0, 1)$ is a small positive constant. For computing the step size α_n, we propose an adaptive local line search algorithm, which is inspired in the backtracking line search scheme used in the Quasi-Gauss-Newton algorithm with an energy reduction constraint [29]. ANLCG algorithm (Step 2) requires that α_n is accepted only if it guarantees a sufficient reduction in the energy $U(f)$, that is, if

$$U(f_n + \alpha_n s_n) \leq (1 - \varepsilon)\, U(f_n),$$

to achieve this we do the following:

Computation of α_n (Backtracking line search algorithm with inertia)

Initially set $a = 0.01, m = 0$, and δ small enough (e.g. $\delta = 10^{-4}$).

2.1 $\alpha_n = a$

2.2 While $U(f_n + \alpha_n s_n) > (1 - \alpha_n \delta)\, U(f_n)$
$\alpha_n = \alpha_n / c_1$, $m = 0$

2.3 $m = m + 1$

2.4 if $m > c_2$
$a = c_3 \alpha_n$, $m = 0$
else
$a = \alpha_n$

This algorithm may be understood as follows: if the computed step length, α_n, decreases sufficiently the energy, then α_n is accepted; otherwise, the step length is reduced by a factor of $1/c_1$. The inertia is introduced by counting the number of consecutive iterations, m, such that the same value of α was accepted. The inertia prevents α being too small. As the method performs an inexact line search (the optimum α is not computed in each iteration), β (step 5 in ALNCG) is computed using the Polak-Ribière formula with restart.

Other algorithms for computing α (with the minimum size requirement), are based on the iterative estimation of the minimum of a quadratic interpolation along the search direction. However, this involves the additional computational cost of evaluating the energy, or the gradient [29] [30]. On the other hand, the Gauss-Newton algorithm requires to compute the product of the approximated Hessian and the gradient vector [29], which is also computational expensive. In the case of the PARC-IL, we observe that although the number of iterations are reduced with the use of better line search algorithms, the resulted computational time is larger than the one obtained using the strategy proposed here.

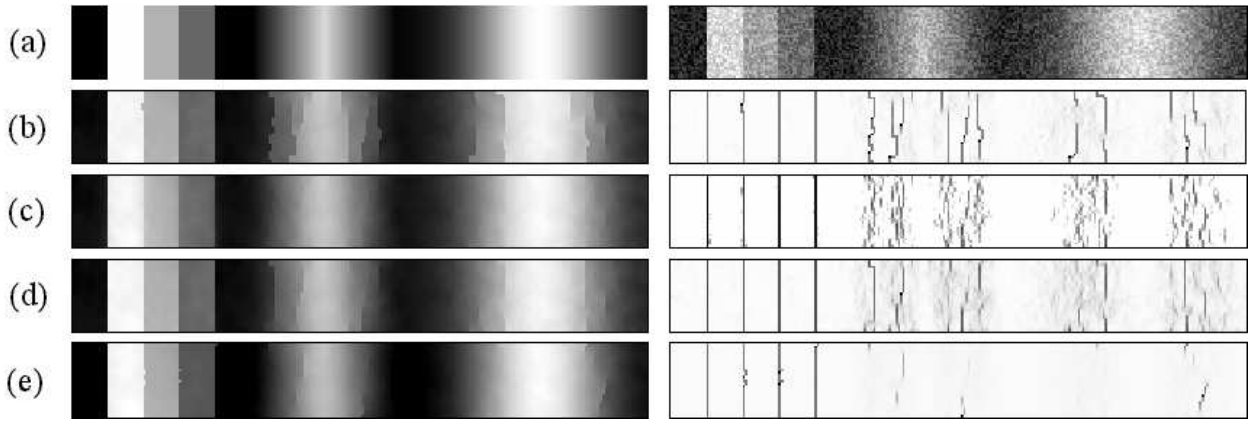

Figure 4: (a) Synthetic real image (left column) and corrupted noisy data test image (right column). (b)-(e) Reconstructions computed with the first order models, potentials: Geman-McClure, Huber (convex), Teboul and SARC-G, respectively. The corresponding edge variables are shown in the right column.

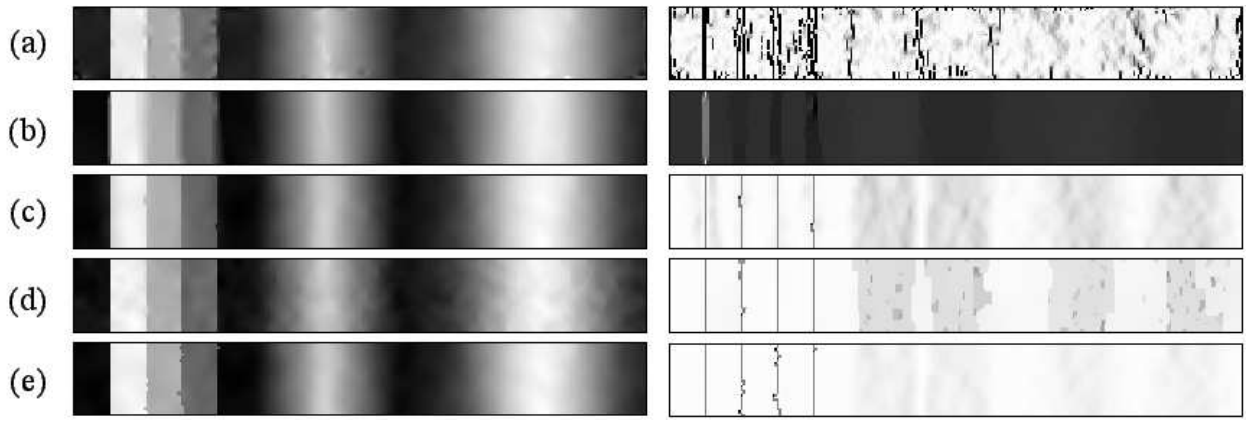

Figure 5: Computed reconstructions and auxiliary variables corresponding to the data in panel 4-(a). (a) Weak plate potential. (b) Proesmans potential. (c) PARC-EL. (d) PARC-EL$^+$. (e) PARC-IL. The corresponding edge variables are shown in the right column.

Empirically, we have found that the values of parameters $c_1 = 2$, $c_2 = 5$ and $c_3 = 3$ and the initial step size a work properly. Note that since α_n ensures that the energy decreases at every iteration, the convergence and stability of algorithm ANLCG are automatically guaranteed.

5 Experiments

In this section we present the results of experiments (with both synthetic and real data) that demonstrate the performance of the ARC–based methods presented above, and compare it with that of classical, regularization–based EPR techniques.

5.1 Comparison with the half-quadratic second order model

The first experiment is a comparison of the performance of the SARC-G models with respect to other first order EPR models: in figure 4, row (a) we show the synthetic test image (first column) and the noise

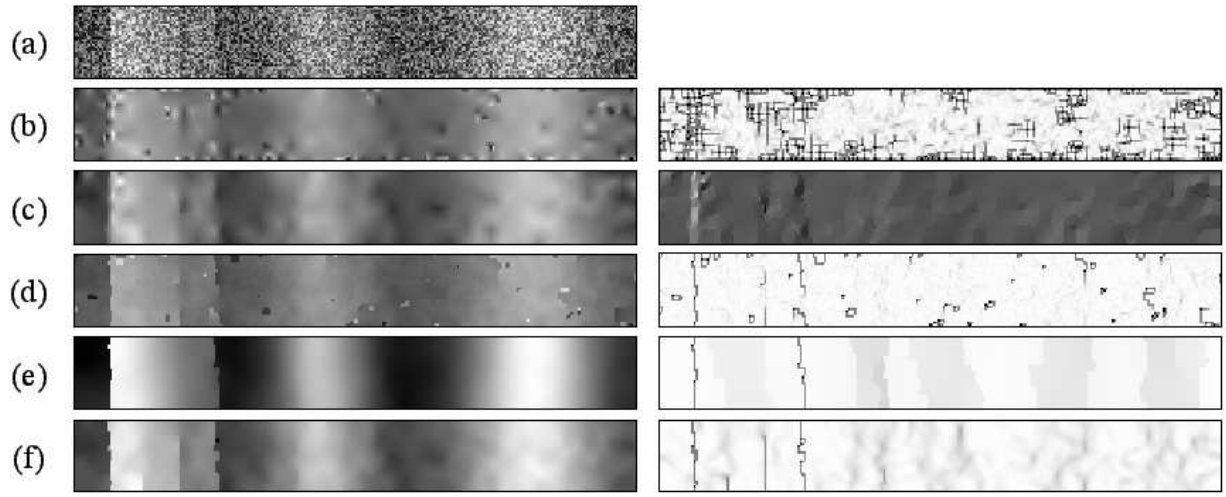

Figure 6: Filtering of Low SNR images. (a) Corrupted with uniform noise with amplitude equal to 100% of the dynamic range. (b) Weak plate potential. (c) Proesmans potential. (d) PARC-EL. (e) PARC-EL$^+$. (f) PARC-IL. The corresponding edge variables are shown in the right column.

corrupted image (second column). The test image was designed so that it includes the following kind of regions: piecewise constant (delimited by first order discontinuities); piecewise constant slope regions (delimited by second order discontinuities) and smooth (non-planar) regions. The data were normalized to the interval [0,1] and corrupted with uniform noise with amplitude equal to 0.2, the MSE for the noisy data is 1.33×10^{-2}. The image dimension is 32×256 pixels. Row (b), shows the reconstruction computed with the WSS model. The potential function used corresponds to the one reported in [31]. This potential has been extensively used and has demonstrated its superior performance with respect to other first order EPR potentials (see [13] [16]). As one can see, in spite of the fact that the regions with constant and smooth changes in the slope are reconstructed with an acceptable quality, the gray level steps are over-smoothed. The results of row (c) correspond to the convex potential better evaluated in [12]. This potential corresponds to the Huber's M-Estimator [6][32]. We can note that convex potentials have a tendency to reduce the dynamic range of the jumps producing a blurring effect[12]. Row (d) shows the results computed with the potential reported by Teboul et al. [26]. This algorithm improves the quality of the first order discontinuities by robustly penalizing the gradient of the auxiliary variables. Finally, results in row (e) correspond to the SARC-G model. The MSE for the restoration in rows (b)-(e) are 9.6×10^{-4}, 6.5×10^{-4}, 5.1×10^{-4} and 4.7×10^{-4}, respectively. One can note that first order EPR techniques introduce edges at the regions with high slope, this effect being observed also with the SARC-G potentials.

Figure 5 shows the restorations and auxiliary variables computed with second order potentials for the data used in figure 4. Row (a) shows the results for the second order model reported in [5]. A convex version of this second order potential is reported in [33]. Images in row (b) corresponds to the results using the potential corresponding to equations 9.14 in [24]. One can appreciate that in both cases, the reconstructed images present ramps in first order discontinuities, where there should be jumps. Other formulations for anisotropic diffusion presented in [24] introduce overshoots at the edges and promote edges in the ramps with high gradient. Rows (c) (d) and (e) show the reconstructions computed with the proposed PARC-EL, PARC-EL$^+$ and PARC-IL models, respectively. The auxiliary variable l in PARC-EL$^+$ shows clearly high order discontinuities between regions with almost constant slope. The MSE for the restorations in rows (a)-(e) are 6.65×10^{-4}, 14.9×10^{-4}, 5.01×10^{-4}, 4.8×10^{-4} and 4.1×10^{-4}, respectively. PARC potentials are slightly winners in terms of the MSE, but they are the clear winners in a visual inspection of the results. As one can note, PARC potentials reconstruct the three different kinds of regions with high quality.

In our experiments, PARC models also have shown better performance for low signal to noise ratios.

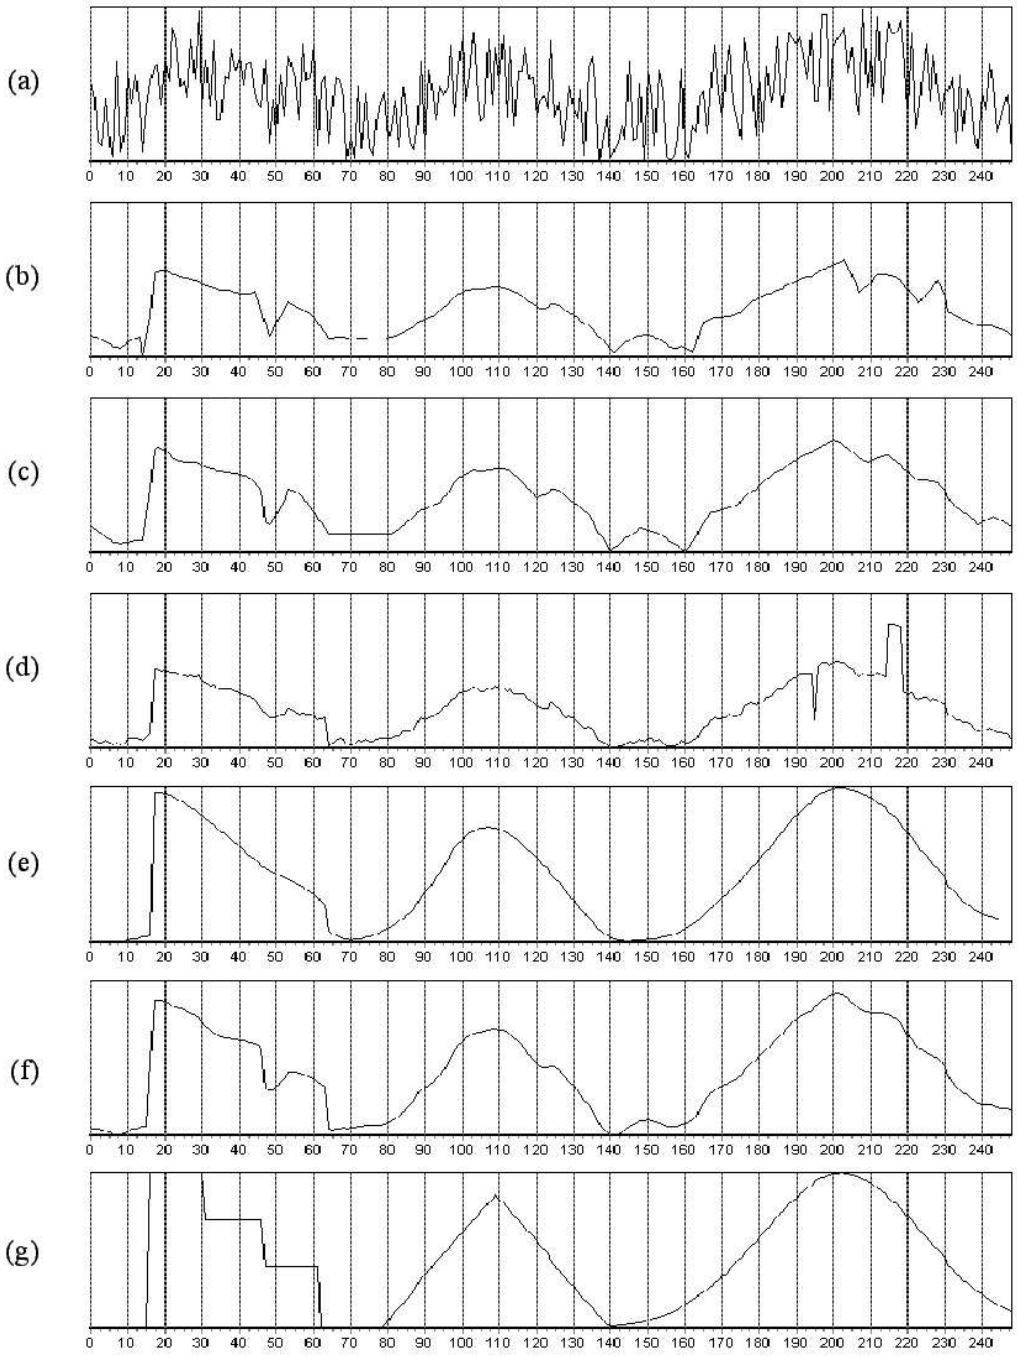

Figure 7: Rows (a)-(f): plots of the center rows of the images in left column in Fig. 6. (g) Ground truth. The profiles were scaled for display purposes.

Figure 6 shows the computed restorations for very noisy data. Row (a) shows the corrupted image with uniform noise with amplitude equal to the dynamic range of the data. Row (b) shows the restored and auxiliary variable computed with the weak thin plate model (WTP) (same method than in row (a) of figure 5). Row (c) shows the result computed with the potential reported by Proesmans et al. [24]. Rows (d), (e) and (f) show the results computed with PARC-EL, PARC-EL$^+$ and PARC-IL, respectively. The MSE for the reconstructions in the rows 6-(a) to 6-(f) are sumarized in table 1. Figure 7 shows the profiles corresponding to the center row for the images in Fig. 6. From these experiments one can conclude that, for images corrupted with moderated noise, the PARC-IL model produces sharper restorations than PARC-EL or PARC-EL$^+$. However, for low SNR, the restorations with PARC-EL$^+$ are better. In subsection 5.4, we discuss the criteria for choosing the parameters of the PARC models.

Table 1. Summary of the results corresponding to figure 6.

ROW	METHOD	MSE ($\times 10^{-4}$)	Parameters
6-(b)	WTP	321.28	$\lambda = 50$, $\mu = 1/25$
6-(c)	Proesmans	157.21	$\lambda = 50$, $\mu = 1/50$, $k = 1/1e6$
6-(d)	PARC-EL	99.01	$\lambda = 1000$, $\mu = 1/50$
6-(e)	PARC-EL$^+$	35.75	$\lambda = 1000$, $\mu = 1/300$, $\gamma = 0.1$, $k_l = 1/50$
6-(f)	PARC-IL	48.56	$\lambda = 1000$, $k = 1/100$, $\varepsilon = 0.05$

5.2 Restorations of real Images

In order to illustrate the performance of the second order models, we performed the following experiment with a real test image. In Figure 8, panel (a) is shown the cameraman picture (256×256 pixels) corrupted with Gaussian noise with $\sigma = 20\%$ of the dynamic range of the original image (which was previously normalized to $[0, 1]$). Rows (b) to (f) show the computed filtered images and edges. The summary of these results is presented in table 2. Regardless of the MSE, one can see that the results computed with the PARC models present best restorations and sharp edges. In this experiment, all the methods were implemented using only vertical and diagonal cliques. The used computer was a pentium III at 800 MHz. See subsection 5.4 for a discussion about the criteria for choosing the parameters of the PARC models.

Table 2. Summary of the results corresponding to figure 8.

ROW	METHOD	MSE ($\times 10^{-4}$)	Time (secs.)	Parameters
8-(b)	WTP	44.72	190	$\lambda = 10$, $\mu = 1/20$
8-(c)	Proesmans	42.33	80	$\lambda = 10$, $\mu = 1/50$, $k = 1/1e6$
8-(d)	PARC-EL	40.50	62	$\lambda = 4$, $\mu = 1/35$
8-(e)	PARC-EL$^+$	43.21	82	$\lambda = 4$, $\mu = 1/30$, $\gamma = 0.1$, $k_l = 50$
8-(f)	PARC-IL	43.33	124	$\lambda = 70$, $k = 1/100$, $\varepsilon = 0.05$

5.3 Anisotropic Diffusion based on PARC potentials

In many cases, the best value for the regularization parameter λ in Eq. (14) is not known in advance, and must be determined in a trial and error basis. In these cases it is often better to use a generalized anisotropic diffusion scheme, as in Refs. [24][27] [34], so that the time itself (i.e., the iteration number) acts as a regularization parameter, which may then be fixed interactively. A PARC–based anisotropic diffusion takes then the form:

$$f_r^{t+1} = f_r^t + \alpha \frac{\partial}{\partial f_r} U_d\left(f^t\right) \tag{15}$$

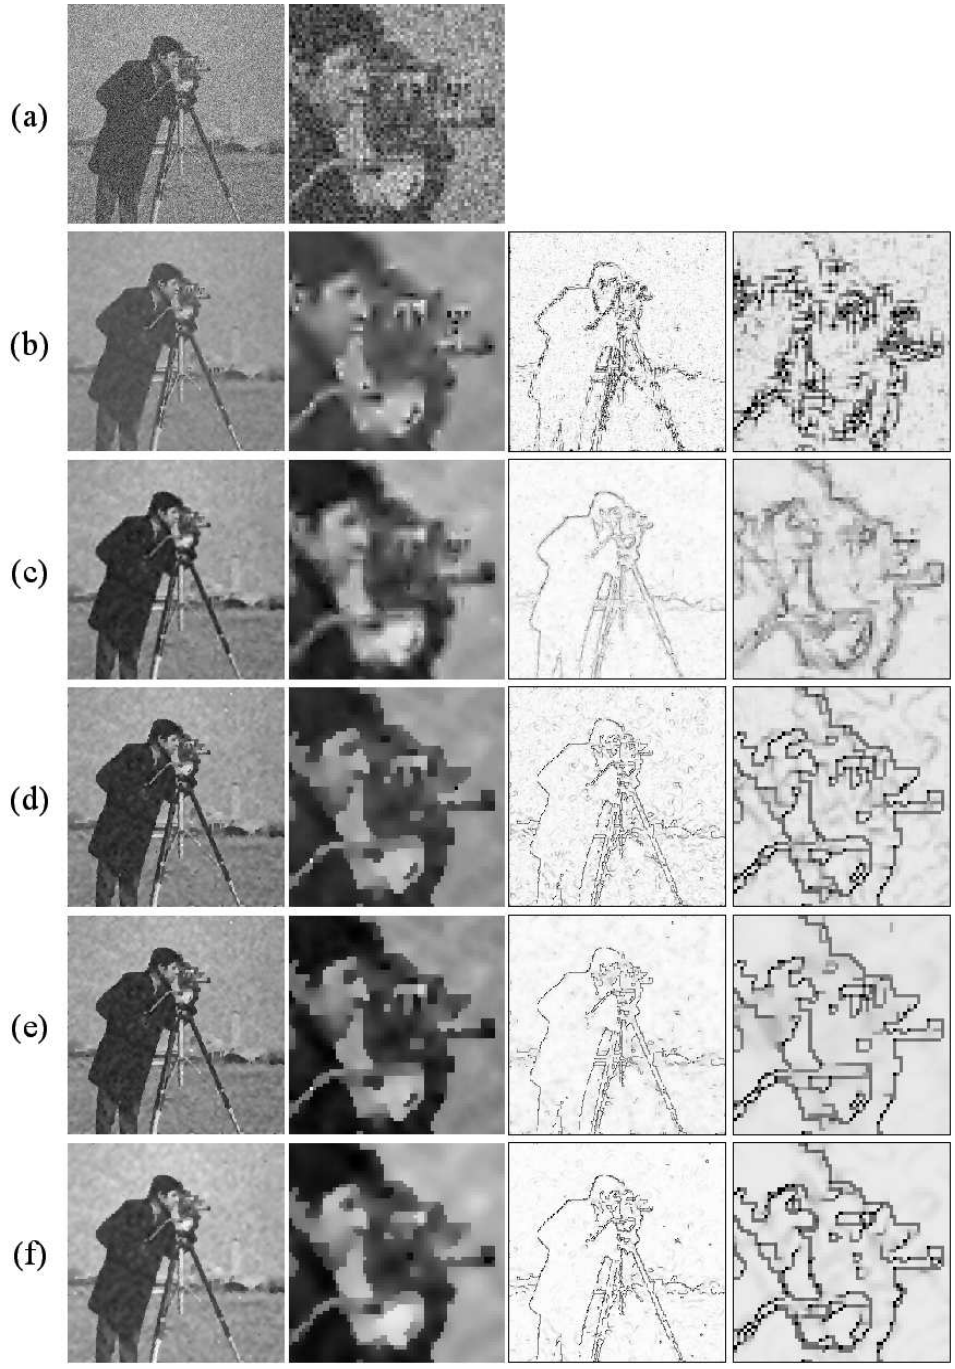

Figure 8: (a) Cameraman picture corrupted with gaussian noise $\sigma = 20\%$. Reconstructions computed with (b) WTP, (c) Proesmans potential, (d) PARC-EL, (e) PARC-EL$^+$ and (f) PARC-IL. The 2 rightmost columns show the corresponding edge maps.

Figure 9: Images of a sequence computed with an anisotropic diffusion algorithm based on PARC–IL potentials

where U_d is computed from (14), dropping the data term. In the PARC-IL case, we get:

$$U_d(f,l) = \sum_{q,s \in N_r} |\Delta^+ f_r - \phi(\Delta^+ f_r) - (\Delta^- f_r - \phi(\Delta^- f_r))|^2 \tag{16}$$

with ϕ given by (9). The step size α in (15) is chosen such that:

$$U_d\left(f^{t+1}\right) \leq U_d\left(f^t\right).$$

Thus, we compute α in at each iteration using the backtracking algorithm with inertia presented in section 4.2. Figure 9 show images of a sequence computed with this method. f^0 corresponds to the cameraman picture corrupted with Gaussian noise with $\sigma = 5\%$.

5.4 Parameter Selection of the Models PARC

In the experiments presented, the parameters for each method were empirically selected, so that in each case the best result was obtained. In the experiments with the synthetic images (figures 4, 5 and 6), we chose a large value for the smoothness parameter, λ, given our prior knowledge about the smoothness of the solution. This smoothness parameter was chosen small for the cameraman picture because we expect a larger variation in the regions surrounded by edges. The μ parameter penalizes the detection of edges; therefore, μ is larger with respect to λ in the synthetic image than in the cameraman, where more edges are expected. In the case of PARC-EL$^+$, the μ parameter is chosen slightly larger than in PARC-EL, in order to compensate for the hysteresis effect of the line potentials.

6 Conclusions

We have presented a new model for edge-preserving regularization. This model is based on the physical analogy of adaptive rest condition potentials (ARC), which are given by the product of a function ϕ of the image gradient times an edge indicator variable l. In order to focus the paper on the characteristics of the ARC potentials, we dealt with the problem of restoring noisy images, but their use can be extended to other image processing and computer vision problems as well. We showed that, for particular choices of ϕ and l, the first order ARC model is equivalent to the well known adaptive weak spring model, but in the general case, one gets a new algorithm that exhibits a better behavior, and, more importantly, that can naturally be extended to generate second (and higher) order potentials; we called these potentials the thin plate model

with adaptive rest condition (PARC). These PARC potentials have the property of automatically changing their behavior to a first order EPR potential at the edges (steps in the gray level) of the image. As a result, regularized cost functionals based on the PARC model are more stable and perform a better restoration of edges and smooth regions, because PARC potentials extend the definition of smoothness to include regions with almost constant slope.
We introduced two kinds of PARC potentials: the PARC with explicit line process (PARC-EL) and the PARC with implicit line process (PARC-IL). We found that these families have specific advantages:
PARC-EL potentials generate a pair of coupled systems (in general, linear for the restored image and non-linear for the auxiliary variables) that can be alternatively minimized. A special case results from selecting a coupled quadratic potential; in such case the resulting coupled systems are linear and can be efficiently minimized (in our case, we used the Gauss-Seidel algorithm in an alternated scheme). This model can incorporate potentials that penalize specific configurations of the auxiliary variable (for example the thickness of the edges; see [13, 26, 27] for more details).
The experiments have shown that (in general) one obtains better reconstructions with the PARC-IL model although at a higher computational cost. However, if the line enhanced process is applied (PARC-EL$^+$) the algorithm becomes more robust when processing images with low SNR.

An interesting open theoretical problem is to determine the precise relation between PARC-EL and PARC-IL potentials, as has been found in the case of WSS models.

7 Acknowledgments

This work was supported in parts by a grant from Conacyt, Mexico (34575-A). The authors also thank to the anonymous referees, since their comments helped to improve the quality of the paper.

References

[1] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and Bayesian restoration of images," IEEE Trans. Pattern Anal. Machine Intell., 6, 1984, 721–741.

[2] D. Terzopoulos, "Regularization of inverse visual problems involving discontinuities," IEEE Trans. Pattern Anal. Machine Intell., 8, 1986, 413–424.

[3] J. L. Marroquin, S. Mitter and T. Poggio, "Probabilistic solitions of ill–posed problems in computer vision," J. Amer. Stat. Assoc., 82, 1987, 76–89.

[4] A. Blake and A. Zisserman, Visual reconstruction, The MIT Press, Cambridge, Massachusetts, 1987.

[5] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," IEEE Trans. Image Processing, 14, 1992, 367–383.

[6] D. Shulman and J. Y. Herve, "Regularization of discontinuous flow fields," Proc. Workshop Visual Motion, 1989, 81-86.

[7] R. L. Stevenson, B. E. Schimitz and E. J. Delp, "Discontinuity preserving regularization of inverse visual problems," IEEE Trans. on SMC, 24, 94, 455–469.

[8] C. Schnörr, "Unique reconstruction of piece wise smooth images by minimizing strictly convex non-quadratic functionals," J. Math. Imag. Vis., 4, 1994, 189–198.

[9] T. Chan, P. Blomgrem, P. Mulet and C. K. Wong, "Total variation image restoration: Numerical methods and extensions," in Int. Conf. Image Processing'97, vol. III, 1997, 384–387.

[10] J. Heers, C. Schnörr and S. Stiehl, "Globally convergent iterative numerical schemes for nonlinear variational image smoothing and segmentation on a multiprocessor machine," IEEE Trans. Image Processing, 6, 2001, 852–864.

[11] J. Idier "Convex half-quadratic criteria and interacting auxiliary variables for image restoration," IEEE Trans. Image Processing, 10, 2001 1001–1009..

[12] S. Z. Li, "Close-form solution and parameter selection for convex minimization-based edge-preserving smoothing," IEEE Trans. Pattern Anal. Machine Intell., 20, 1998, 916–932.

[13] M. J. Black and A. Rangarajan, "Unification of line process, outlier rejection, and robust statistics with application in early vision," Int. Journal of Computer Vision, 19, 1996, 57–91.

[14] D. Geman and C. Yang, "Nonlinear image recovery with half-quadratic regularization," IEEE Trans. Image Processing, 4, 1995, 932–946.

[15] L. Cohen, "Auxiliary variables and two-steps iterative algorithms in computer vision problems," Journal Mathematical Imaging and Vision, 6, 1996, 59–83.

[16] P. Charbonnier, L. Blanc-Féraud, G. Aubert and M. Barlaud, "Deterministic edge-preserving regularization in computer imaging," IEEE Trans. Image Processing, 6, 1997, 298–311.

[17] T. Kubota and T. Huntsberg, "Adaptive anisotropic parameter estimation in the weak membrane model," Proc. Int WS in EMMCVPR'97, Lecture Notes in Computer Vision 1223, Springer Verlag, Venice Italy, 1997, 179–194.

[18] A. Chambolle and P. L. Lions, "Image recovery via total variation minimization and related problems," Numer. Math. 76, 1997, 167–188.

[19] L. Ambrosio and V. M. Tortorelli, "Approximation of functionals depending on jumps by elliptic functionals via Γ-convergence, " Commun. Pure Appl. Math, 43, 1990, 999–1036.

[20] Y.-L. You and M. Kaveh, "Blind image restoration by anisotropic regularization," IEEE Trans. Image Processing, 8, 1999, 396–407.

[21] M. Rivera and J. L Marroquin, "The adaptive rest condition spring model: an edge preserving regularization technique," in Proc. IEEE-ICIP 2000, Vol. II, Vancouver, BC, Canada, 2000, 805–807.

[22] M. Rivera, and J. L. Marroquin, "Efficient Half-Quadratic Regularization with Granularity Control," Technical Report: 05.06.2001, I-01-07 (CC), 2001, CIMAT, Mexico.

[23] M. Rivera, and J. L. Marroquin, "Adaptive Rest condition Potentials: Second Order Edge–Preserving Regularization," Proc. European Conf. Computer Vision 2002, ECCV'02, A. Heyden, G. Sparr, M. Nielsen and P. Johansen, editors, LNCS 2350, Copenhagen Denmark, May 2002, 113–127.

[24] M. Proesmans, A. E. Pouwels and L. Van Gool, "Couple geometry-driven diffusion equations for low level vision," In B. M. ter Haar Romeny, editor, Geometry-Driven Diffusion in Computer Vision, 191–228, Kluwer Academic Publisher B.V. 1994.

[25] M. Nitzberg and T. Shiota, "Nonlinear image filtering with edge and corner enhancement," IEEE Trans. Pattern Anal. Machine Intell., 14, 1992, 826–833.

[26] S. Teboul, L. Blanc-Féraud, G. Aubert and M. Barlaud, "Variational approach for edge-preserving regularization using coupled PDE's," IEEE Trans. Image Processing, 7, 1998, 387–397.

[27] M. J. Black, G. Sapiro, D. H. Marimont and D. Heeger, "Robust anisotropic diffusion," IEEE Trans. Image Processing, 7, 1998, 421–432.

[28] S. Z. Li, Markov Random Field Modeling in Image Analysis, Springer-Verlag, Tokyo, 2001.

[29] J. Nocedal and S. J. Wright, Numerical Optimization (Springer Series in Operations Research), Springer-Verlag, New York, 1999.

[30] J. J. Moré and D. J. Thuente, "Line search algorithms with guaranteed sufficient decrease," ACM Trans. on Mathematical Software, 20, 1994, 286–307.

[31] S. Geman and D. E. McClure, "Bayesian image analysis methods: An application to single photon emission tomography," in Proc. Statistical Computation Section, Amer. Statistical Assoc., Washington, DC, 1985, 12–18.

[32] P. J. Huber, "Robust Statistics," Wiley, New York, 1981.

[33] O. Scherzer, "Denoising with high order derivatives of bounded variation and an application to parameter estimation," Computing, 60, 1998, 1–27.

[34] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," IEEE Trans. Pattern Anal. Machine Intell., 12, 1990, 629–639