Efficient Half–Quadratic Regularization with Granularity Control

Mariano Rivera and Jose L. Marroquin
Centro de Investigacion en Matematicas A.C.
Apdo. Postal 402, Guanajuato, Gto. Mexico 36020
email: mrivera@cimat.mx
Telephone +52 (473) 73 271 55
Fax: +52 (473) 73 257 49

Abstract

In the last decade, several authors have proposed edge preserving regularization (EPR) methods for solving ill posed problems in early vision. These techniques are based on potentials derived from robust M–Estimators. They are capable of detecting outliers in the data, finding the significant borders in noisy images and performing edge–preserving restorations. These methods, however, have some problems: they are computationally expensive, and often produce solutions which are either too smooth or too granular (with borders around small regions). In this paper we present a new class of potentials that permits to separate robustness and granularity control, producing better results than the classical ones in both scalar and vector-valued images. We also present a new fast, memory-limited minimization algorithm.

Keywords. Half–Quadratic Regularization, Edge–Preserving Regularization, Image Restoration, Non-linear Filtering, Conjugated Gradient Algorithms.

1 INTRODUCTION

In the fields of image processing, image analysis and computer vision, one deals with the problem of reconstructing an image $\widehat{f}$ from noisy and degraded observations g. Consider the following model of the observations:

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

where η is additive noise and F is a linear operator that is assumed to be known. For example in optical blurring, F corresponds to the convolution with the Point Spread Function (PSF) of the imaging system.

In other cases, F can be a linear approximation of a non-linear transformation (e.g. in the computation of optical flow).

The information provided by the data and the direct model (1) is, in general, not enough for an accurate estimation of $\widehat{f}$, so that the regularization of the problem is necessary. That means that, *a priori* information or assumptions about the structure of $\widehat{f}$ need to be introduced in the reconstruction process. The regularized solution f^* is computed by minimizing an energy functional U:

$$f^* = \arg\min_f U(f) \tag{2}$$

where

$$U(f) = D(f) + \lambda R(f). \tag{3}$$

The data term D establishes that the reconstruction f should be consistent with the data g. The regularization term R imposes a penalty for violating the *a priori* assumptions, and the relative contribution of each term to the global energy is weighted by the parameter λ. The Data term can been written as

$$D(f) = \sum_r \rho_D \left(t_r \left(f\right)\right)$$

where t is the residual error defined by

$$t_r \left(f\right) = F(f)_r - g_r,$$

where $r = (x, y)$ represents a site in the pixel lattice L, and ρ_D is a potential function that defines the error norm; the subindex denotes that this potential is associated with the data term. In the framework of Bayesian regularization [1], D is chosen as the negative log–likelihood and the prior constraints are incorporated in the form of *a priori* (Markov Random Filed) model for f, so that $R(f)$ 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, so that $R(f)$ takes the form:

$$R(f) = \sum_r \sum_{s \in N_r} \rho_R(u_{rs} \left(f\right)).$$

where the residual error for the regularization term is defined by

$$u_{rs} \left(f\right) = f_r - f_s. \tag{4}$$

and ρ_R is a potential function. For instance, the homogeneous spring model R_S is obtained by assuming that η corresponds to Gaussian noise:

$$\rho_D \left(t_r \left(f\right)\right) = \left(F(f)_r - g_r\right)^2. \tag{5}$$

and choosing ρ_R as a quadratic potential over the first differences:

$$\rho_R(u_{rs}(f)) = d_{rs}(f_r - f_s)^2 \tag{6}$$

where d_{rs} is a constant that depends on the distance between pixels r and s, for instance: $d_{rs} = |r-s|^{-1}$. This quadratic potential corresponds to the *a priori* assumption that the original data $\widehat{f}$ is globally smooth. Then, assuming that F is linear, the cost functional that results from potentials (5) and (6) is quadratic:

$$U_1(f) = \sum_r \left\{ |t_r(f)|^2 + \lambda \sum_{s \in N_r} d_{rs} |u_{rs}(f)|^2 \right\} \tag{7}$$

This cost functional is not robust to outliers. In the image restoration context, the outliers are located at those sites where the assumptions implicit in the cost function are not fulfilled. In particular, for a regularization term that assumes global smoothness, the outliers correspond to the edges in the image. As a consequence, the potential function (6) will produce an over–smoothing of the real edges of the image.

To alleviate this problem, there have been proposed robust potential functions for the data and regularization terms. This regularization technique is usually based on potentials derived from robust M–Estimators and is capable of detecting outliers in the data, finding the significant edges of a noisy image and performing an edge–preserving restoration. However, there are some problems: the robust potentials that are in use have a single parameter that controls the minimum residual magnitude that corresponds to an outlier; it often happens in noisy images that if this parameter is too small, many small regions generate edges around them so that the solutions appear granular, while if the value of this parameter is increased, some true edges are not preserved Another problem is that the convergence of the algorithms that have been proposed for the minimization of the corresponding cost function are relatively slow, making these methods computationally expensive.

The purpose of this paper is to present a formulation for robust potentials that produces faster and better behaved algorithms that result in better reconstructions and a significant time processing reduction.

The organization of the paper is as follows: in section 2, a brief introduction to robust regularization is presented. Section 3 introduces the new formulation for robust potentials. Also, in that section, we show how to write robust energy terms for vector valued data. In section 4, we present a non-linear Conjugated Gradient algorithm for half–quadratic regularized functionals with minimal memory requirements. Experiments that demonstrate the performance of the new potentials and the algorithm introduced herein are presented in section 5. In section 6, we present a heuristic method for selecting good values for the parameters of the method. Finally, our conclusions are given in section 7.

2 REVIEW OF ROBUST REGULARIZATION BASED ON M-ESTIMATORS

To solve the problem of over–smoothing the real borders in f, potential functions for the data and regularization terms that increase at a smaller rate than the quadratic potential have been used; see Refs. [3][4][5] [6][7][8][12] [10] and first order ARC potentials in [11]. These robust potentials $\rho(t)$ (that correspond to robust objective functions used in the statistical literature[7],) have the following characteristics [8][13][14]:

1. $\rho(t) \geq 0 \ \forall t$ with $\rho(0) = 0$,
2. $\psi(t) \equiv \partial \rho(t) / \partial t$ exists.
3. $\rho(t) = \rho(-t)$, the potential is symmetric
4. $\omega(t) \equiv \frac{\psi(t)}{2t}$ exists
5. $\lim_{t \to +\infty} \omega(t) = 0$
6. $\lim_{t \to 0^+} \omega(t) = M$, $0 < M < +\infty$.

(8)

Condition (8.1) establishes that a residual equal to zero must produce the minimum cost. Condition (8.2) constrains ρ to be differentiable, so that one can use efficient deterministic algorithms for minimizing the cost function (really, to compute a local minimum). Condition (8.3) constrains ρ to penalize equally positive and negative values of t. Finally, conditions (8.4 to 8.6) imposes the robustness condition. A robust potential corresponds to (in general) a non-convex potential that grows at a slower rate than the quadratic one. A local minimum of the robust–potential–based energy functional

$$U_2(f) = \sum_r \left\{ \rho_D\left(t_r(f)\right) + \lambda \sum_{s \in N_r} \rho_R(u_{rs}(f)) \right\} \tag{9}$$

is computed by solving the system

$$\psi_D\left(t_r(f)\right) + \lambda \sum_{s \in N_r} \psi_R(u_{rs}(f)) = 0. \tag{10}$$

where the derivatives

$$\psi_D(t_r(f)) = \partial \rho_D(t_r(f)) / \partial f_r \quad \text{and} \quad \psi_R(u_{rs}(f)) = \partial \rho_R(u_{rs}(f)) / \partial f_r$$

are called the *influence functions*. One can also define:

$$\omega_r = \frac{\psi_D(t_r(f))}{2t_r(f)} \quad \text{and} \quad \omega_{rs} = \frac{\psi_R(u_{rs}(f))}{2u_{rs}(f)} \tag{11}$$

as the *weight functions*. Then (as was shown in [4][5][8][13]), one can solve the non–linear system (10) with the following two–step iterative algorithm:

Algorithm 1. Weighted Linear System

Given an initial guess for f,

Repeat:

1. Compute the weights ω_r and ω_{rs} using (11).
2. Solve the system

$$\omega_r t_r(f) + \sum_{s \in N_r} \omega_{rs} u_{rs}(f) = 0$$

for f, keeping ω_r and ω_{rs} fixed.

Until convergence.

(12)

The ARTUR algorithm reported in [8] corresponds to algorithm 1 but using as starting point for f an homogeneous image equal to zero.

In References [7] and [8] one can find pictorial summaries of the robust potential function used in robust regularization. One may define a classification of potentials functions based on the shape of the influence functions ψ [14]:

1. Monotone ψ (MT). $\psi(t)$ is constant for $|t| \geq k$. Where k is a given threshold An example is the Huber's potential function [13]:

$$\rho_h(t;k) = \begin{cases} t^2 & |t| \leq k \\ k(2|x| - k) & otherwise. \end{cases} \tag{13}$$

2. Soft redescender ψ (SR). One has $\psi(\infty) = 0$. An example is the Cauchy's potential function:

$$\rho_c(t;k) = \log(1 + (t/k)^2),$$

where k is the scale parameter.

3. Hard redescender ψ (HR). One has $\psi(t) \approx 0$ for $|t| \geq \theta$, where θ is a given threshold. This kind includes the widely used Welsh's potential function:

$$\rho_w(x,k) = 1 - \frac{1}{2k}\exp(-kx^2), \tag{14}$$

and Tukey's (biweight) potential function:

$$\rho_b(t;k) = \begin{cases} (1 - (1 - (2t/k)^2)^3) & |t/k| < 1/2 \\ 1 & otherwise; \end{cases}$$

where k is the scale parameter.

The corresponding influence functions and the weight functions are also illustrated. In order to choose the "right" ρ–function for computer vision problems, we need take into account the following:

1. Robustness to outliers. That is, one wants to reduce the effect of large errors in the solution.

2. Computational difficulties. The minimization algorithm must be stable and have fast convergence.

In the robust statistics literature [17] one can find measures of robustness for potential functions. These measures are derived from the influence function; they show the influence of large errors (outliers) and small errors (measurement errors); other common measure is the *breakdown point* (that in general does not depend on the shape of the potential ρ) that measures the robustness with respect to a large quantity of outliers. The HR (potentials that reject extreme outliers completely) and SR (potentials that reduce considerably the effect of extreme outliers) classes are preferred for regression and estimation problems in both contexts: statistics [14][16][17] and image processing–computer vision [3][4][6] [7][8][10][18] [19][20][21][22].

For the case that concerns this paper (image restoration), the robust methods are capable of detecting the significant outliers in the data, finding edges (outliers with respect the *prior* assumption) of a noisy image and performing an edge–preserving restoration. However, the HR–potentials do not guarantee uniqueness of the solution [7]. As a consequence, a good initial guess must be provided in order to avoid to be trapped by a "bad" local minimum. Another problem is that there is no explicit control of the granularity of the solution, and as a result one may have either inaccurate definition of edges or an over detection of small details (see Figure 2). Finally, the convergence rate of the minimization algorithm 1 is relatively slow, so that it is computationally demanding.

In an attempt to solve these problems, in [8] the initial point is chosen as an homogeneous image equal to zero (algorithm ARTUR). This reduces the granularity, but it also reduces the accuracy of the computed edges (see 2–f) and does not contribute to the reduction of the computation time. Another improvement, introduced in [7], is to make the potential in the data term robust, which allows one to eliminate unstructured outliers; this however, may cause Algorithm 1 to become unstable because the weights ω_r and ω_{rs} can both be equal to zero (because of condition (8.5)), or produce ill–conditioned systems for small values of these weights. Also, this formulation does not consider that in the image processing context, outliers in the data may be structured (there may be small and well defined noisy regions, see Figure 2–c), so that it is not clear how to constrain granularity. In [7][9][10] and [19] potentials that penalize thickness and promote the continuity of the borders are introduced. The result is a better definition of the edges, but the computational time is increased because there is no closed formula for the weights, so that they must be computed for each iteration as the solution of a non–linear system which causes a significant slow–down of the algorithm.

3 POTENTIALS FOR EDGE PRESERVING REGULARIZATION WITH GRANULARITY CONTROL

In this section we propose a method for stabilizing edge-preserving potentials that produces faster and better behaved algorithms. First, we note that condition (8.5) needs to be redefined in order to avoid that the weights ω_r and ω_{rs} can both be close to zero, ill–conditioning the system (10). The new condition is:

$$\lim_{t \to +\infty} \frac{\rho'_R(t)}{2t} = \mu \tag{15}$$

where $\mu \in (0, 1]$ is a positive parameter. Now, the granularity of the solution has to be controlled, so that a large and well defined region is preferred over a group of small regions (well defined regions are such that the weights at their borders are zero if a HR–potential is used). In order to introduce this control in the robust potential, it is necessary to include an additional term in the cost functional (9), that assigns a small additional cost to large jumps that are assigned a constant cost by the potential ρ_Q. An easy way to stabilize the system and at the same time control granularity is simply to add a quadratic term to the HR–potential ρ_H, so that it has heavier "tails". The resulting regularization potential ρ_Q is given by

$$\begin{aligned} \rho_Q(u_{rs}(f), k, \mu) &= \mu d_{rs}[u_{rs}(f)]^2 \\ &+ (1-\mu) d_{rs} \rho_H(u_{rs}(f), k), \end{aligned} \tag{16}$$

where the μ parameter controls the granularity of the solution and k is a positive scale parameter. Note that for the extreme value $\mu = 1$, ρ_R is quadratic (non-robust) and strongly penalizes the granularity. When μ is close to zero 0, the potential is HR–robust and promotes piecewise smooth reconstructions. Small values for μ allow smooth changes inside the regions and control the size of the grain in the reconstruction. The form of the ρ_Q potential together with its influence and weight functions is presented in Figure 1. This type of influence function corresponds to the so–called Quasi–robust estimators in the statistics literature [23]. In all our experiments we use the Welsh's potential (14):

$$\rho_H(x, k) = \rho_w(x, k),$$

in (16) and for the data term we use a quadratic potential

$$\rho_D(t) = |t|^2,$$

then we have:

$$U_N(f) = \sum_r \left\{ \rho_D(t_r(f)) + \lambda \sum_{s \in N_r} \rho_Q(u_{rs}(f), \mu, k) \right\}, \tag{17}$$

The minimization of (17) with respect to f is computed by solving the non–linear system

$$\frac{\partial}{\partial f_r}\left\{\rho_D\left(t_r\left(f\right)\right)+\lambda\sum_{s\in N_r}\rho_Q\left(u_{rs}\left(f\right),k,\mu\right)\right\}=0$$

this non–linear system can be written in matrix form:

$$W_f f=\widehat{g} \tag{18}$$

where the matrix W_f depends on f and $\widehat{g}_r=[\omega_D]_r g_r$, with ω_D as the corresponding vector of weights for the data term. As ρ_Q satisfies (8) but condition (8.5) is replaced by (15), then, (18) cannot become singular. In section IV, we present an efficient Conjugate Gradient algorithm for solving (18), which has a better convergence rate than Algorithm 1.

3.1 Robust Regularization for Vectorial Data

In the restoration or analysis of vector-valued images $\overline{f}=\left[f_1,f_2,\ldots,f_M\right]^T$ (e.g. optical flow and color image processing) given the data $\overline{g}=\left[g_1,g_2,\ldots,g_M\right]^T$, it is necessarily to couple the process over the M channels.

For illustration proposes, we assume that the channels of $\overline{f}$ are independently acquired, so that the outlier rejection in the data term is decoupled. On the other hand, the outliers corresponding to the regularization term are coupled because we expect that the joint contribution of all channels results on a better detection of the edges. Then the robust potentials for the data and the regularization terms (respectively) are

$$D(\overline{f})=\sum_m\sum_r\rho_D\left(t_r\left(f\right)\right)$$

and

$$R(\overline{f})=\sum_r\sum_{s\in N_r}\rho_Q\left(\left|u_{rs}\left(\overline{f}\right)\right|,k,\mu\right)$$

with $m=1,...,M$, where $\left|u_{rs}\left(\overline{f}\right)\right|^2=\left(f_{1,r}-f_{1,s}\right)^2+\left(f_{2,r}-f_{2,s}\right)^2+\ldots+\left(f_{M,r}-f_{M,s}\right)^2$. The solution $\overline{f}$ is computed by solving

$$\frac{\partial}{\partial f_{m,r}}\left[\rho_D\left(t_r\left(f\right)\right)+\lambda\sum_{s\in N_r}\rho_Q\left(\left|u_{rs}\left(\overline{f}\right)\right|,k,\mu\right)\right]=0. \tag{19}$$

4 MINIMIZATION ALGORITHM

The robust cost functionals presented above are non–quadratic functions with a large number of variables. The minimization algorithms based on the iterative solution of a weighted linear system [e.g., based on Algorithm 1 (10)] as the one used in [4][8] are relatively slow. This is particularly critical in the case of processing three–dimensional (3D) data as in the case of 3D image registration [22]. If one wants to

accelerate the computation of the solution, the use of specialized algorithms for directly solving the non–linear equation system (18) (as the non–linear conjugate gradient (NLCG)[24] or the Newton–type ones[25],) are not a good choice. The reason its that, for example, in order to guarantee convergence, the NLCG algorithms require that at each iteration, the partial solution guarantees a sufficient reduction in the cost $U(f)$ [26], that is

$$U(f_{n+1}) \leq (1-\varepsilon)\, U(f_n).$$

where $\varepsilon \in (0,1)$ is a small positive constant. In order to satisfy this constraint, the NLCG algorithm must perform the expensive evaluation of the cost function at each iteration. On the other hand, the Newton–type algorithms (say, the Gauss–Newton method), additionally need to compute the product of an approximation to the Hessian of $U(f)$ and a vector, this computation being more expensive than evaluating the energy.

This section presents an efficient conjugate gradient algorithm for minimizing (computing a local minimum of) non–quadratic functionals $U(f)$ that are sums of half-quadratic potentials [i.e., potentials that satisfy conditions (8) in its original form or with (15) instead (8.5)]. Before introducing the Half–Quadratic Conjugate Gradient (HQCG) algorithm, we analyze the alternated minimization strategy used by Algorithm 1.

If a potential ρ is half–quadratic, then it has an equivalent representation:

$$\rho(z) = \min_{\omega} \widetilde{\rho}(z,\omega),$$

where $\widetilde{\rho}(z,\omega) = z^2\omega + \Phi(\omega)$ for a given convex potential Φ that penalizes the over detection of outliers [4][5][7] [8]. In the equivalent half–quadratic formulation, one may consider that the weights [given by (11)] result from minimizing the potential $\widetilde{\rho}$ w.r.t. ω. Thus, Algorithm 1 performs an alternated minimization w.r.t. f, ω_r, and ω_{rs}. This means that the non-linear system (18) is solved by the iterative scheme:

$$W_{f_n} f_{n+1} = \widehat{g}_n, \tag{20}$$

where the subindex n denotes the iteration number. In a half-quadratic sense, the computation of the matrix W_{f_n} corresponds the minimization w.r.t. the weights.

The problem with the alternated minimization strategy is that the computation of the weight matrix W and the minimization of the weighted linear system (20) are decoupled, which decreases the efficiency of the method. The method we propose is based instead in the direct minimization of (17) using non–linear conjugated gradient. This general algorithm, however, can be made more efficient, in this case taking advantage of the half–quadratic structure; in particular, this structure allows one to derive a formula for the optimal step size at each iteration: since the minimizer of (15) satisfies (20), where W_{f_n} is a positive definite matrix, the optimal step size $\alpha_n = \min_\alpha U(f_n + \alpha s_n)$ may be computed as

$$\alpha_n = \frac{r_n^T r_n}{s_n^T W_{f_n} s_n},$$

where r_n and s_n are the currents residual vector and descent direction, respectively.

The resulting HQCG algorithm is:

Algorithm 2. HQCG

Set $n = 1, \beta = 0, f_0$ equal to an initial guess.

Repeat:

1. Set $A_n = W_{f_n}$.
2. $r_n = \widehat{g}_n - A_n f_n$
3. if$(n \neq 1)$ $\quad \beta = \max\left\{0, \frac{r_n^T(r_n - r_{n-1})}{r_{n-1}^T r_{n-1}}\right\}$
4. $s_n = r_n + \beta s_{n-1}$
5. $\alpha_n = \frac{r_n^T r_n}{s_n^T A_n s_n}$
6. $f_{n+1} = f_n + \alpha_n s_n$, $n = n + 1$.

Until $|r_n| < \varepsilon$.

Now, we explain each step in detail. Step 1 corresponds to updating the weights ω_r, and ω_{rs}, i.e., one step of the alternated minimization in the half–quadratic sense. The gradient r is updated in step 2. As system A is changing at each iteration, we cannot guarantee that the gradient at iteration n is normal to the last descending direction s_{n-1}(i.e. in general $r_n^T s_{n-1} \neq 0$); so that we compute β with the Fletcher–Reeves (FR) formula with restarting, which is the one generally used in the NLCG algorithms. This FR formula has demonstrated to have better performance in the case of non–linear systems than the Polak–Riviere formula [27].

Note that since we are computing the optimal step size, the condition

$$U(f_n + \alpha_n s_n) < U(f_n).$$

is automatically satisfied.

5 EXPERIMENTS

We present a set of experiments that illustrate the performance and viability of the robust potential with granularity constraint. To isolate the effect of k and μ, for these experiments we use a quadratic potential for ρ_D.

5.1 Segmentation

In the first set of experiments, we study the influence of the 2 parameters that control edge preserving and granularity [i.e. k and μ, respectively in (16)] in the filtered image. The input image (Fig. 2–a) is an axial section of a magnetic resonance image of the brain, with the gray levels normalized in the interval [0,1]. The task in this case is to segment the brain from non–brain tissue, eliminating as much unwanted detail as possible, without distorting the position of the brain /non-brain (B–NB) boundary.

The starting point for HQCG algorithm was the original image. In all the cases we used $\lambda = 200$. Panel 2–b shows the restoration with $k = 80$, and $\mu = 0$ (maximum granularity). As one can see, there is a lot of unwanted detail. If the edge preserving parameter k is decreased (by setting $k = 60$ and $\mu = 0$) the image of the panel 2–c is obtained. Note that some of detail is eliminated (although not completely), but the B–NB boundary is lost. On other hand, if μ is increased to 0.03 (for $k = 800$) one obtains the image of panel 2–d, in which the unwanted detail is smoothed out without dislocating the B–NB boundary. The evaluation of the weight function for the potential $\rho_Q(u_{rs}(f), \mu, k)$ is shown in panel 2–e. As a comparison, in panel 2–f we present the filtered image with the ARTUR algorithm [8], using the Geman–McClude potential $\rho_{GM}(t) = t^2/(k + t^2)$ [3] in the regularization term, which is considered the one that gives the best results [8]. The parameters where hand–adjusted to get the best possible results; their value was $\lambda = 10$, $k = 600$. In this case the starting point was an homogeneous image equal to zero, as is recommended in [8].

5.2 Denoising Vectorial Images

We illustrate the effect of the operation over vectorial images by denoising a color image (channels red (r), green (g) and blue (b)). First row in figure 3 shows the RGB channels of a color image. Each channel is independently corrupted by adding uniform noise with amplitude equal to 60% of the dynamic range of the vector valued image, $\overline{g}$, computed by $\max(\max(g_1), \max(g_2), \max(g_3)) - \min(\min(g_1), \min(g_2), \min(g_3))$. Second row shows the result of processing every channel independently with an HR–potential in the regularization term ($\lambda = 50$, $k = 300$, and $\mu = 0$). As one can appreciate, the edges are not well detected and the images are still noisy. Third row shows the effect of coupling the HR–potential over all the channels ($\lambda = 50$, $k = 75$, and $\mu = 0$). In spite the fact that there is an improvement in the computed images, they are still noisy and the borders are not well preserved. Finally, fourth row shows the computed images that result from coupling the processing over the three channels and considering a penalization over the granularity ($\lambda = 50$, $k = 300$, and $\mu = 0.01$). The improvement is evident: the edges are well preserved and the noise is removed.

5.3 Performance of the HQCG Algorithm

Here we present a comparison of the performance of the algorithms HQCG and ARTUR. The conditions of the test correspond to the recommended in Ref. [8] for the ARTUR algorithm: the starting point was an homogeneous image equal to zero, which contributes to give non–granular results [8], so granularity was not constrained in this experiment; at each iteration both algorithms "introduce" edges. The parameters were $\lambda = 30$, $k = 600$ and $\mu = 0$. The minimization of the half-quadratic cost in ARTUR algorithm was computed by performing 20 iterations of the linear conjugate gradient algorithm (ARTUR–CG).

In figure 4, panel (a) shows the real image used for the test. Panels (b) and (c) show the corresponding filtered images after 30, and 70 iterations (respectively) of the HQCG algorithm. Panel (d) the resulting image after 300 iterations of the ARTUR algorithm. One can note that the dial disc of the telephone is starting to be defined in panel (b) while in the iteration 300 of ARTUR algorithm [panel (d)] is still imperceptible; besides, the edges in the soccer ball are better defined and allocated by HQCG than by ARTUR. The computation times in Fig. 4 correspond a pentium III at 800 Mhz based workstation. The evaluation of the cost function versus the iteration number is plotted in figure 5. As one can see, the HQCG presents a fast convergence rate and a smoother transition between iterations. Note however that one HQCG iteration takes 1.5 times an ARTUR iteration.

5.4 Computation of the Optical Flow

The last experiment illustrates the performance of the presented technique in a classical computer vision problem: the computation of optical flow (OF). Introducing in the Horn-Shunck [28] formulation for OF robust potentials, one gets the energy functional:

$$U(\overline{w}) = \sum_r \left\{ \rho_D \left(t_r \left(\overline{w} \right) \right) + \lambda \sum_{s \in N_r} \widehat{\rho}_Q \left(\left| u_{rs} \left(\overline{w} \right) \right|, k, \mu \right) \right\}$$

where $t_r \left(\overline{w} \right) = \overline{w}_r^T \nabla g_{r,t}$ and $u_{rs} \left(\overline{w} \right) = \overline{w}_r - \overline{w}_s$; the elements of the vector $\nabla g \equiv \left[g_x, g_y, g_t \right]^T$ represent the partial derivatives of the g image with respect to x, y, and t, respectively. The elements u and v of the vector field $\overline{w} = [u, v, 1]^T$ represent the component of the OF in the x and y directions, respectively. The use of standard robust (e.g. HR) potentials for $\widehat{\rho}_Q$ presents, in the case of OF, some particular problems: if there is noise in the data, some small regions may get isolated from the rest of the field by the edges introduced by the algorithm and may then be assigned very large, erroneous displacements. To prevent this undesirable behavior, the granularity control parameter μ becomes specially useful, since it is possible to "anneal" its value by moving it from an initial, relatively large value (e.g. $\mu = 0.3$) to a final small one as the iterations proceed. This results in the final piecewise smooth OF without the undesirable granularity that otherwise may be present. This is illustrated in Fig. 6: panels (a) and (b) show a synthetic image pair with

small square (marked in panel (b)) moving 0.5 pixels down and to the right. Each image is corrupted with Gaussian noise with standard deviation equal to 6% of the dynamic range. Panel (c) shows the magnitude of the OF obtained with the classical Horn-Schunck quadratic potential. The parameter λ was manually set to the value that gets the best result ($\lambda = 5$); panel (d) shows the results obtained with the proposed annealing procedure, where $\mu(t) = 0.3 \times 0.9^{0.02t}$, until the final value of $\mu = 0.04$ is reached. The other parameters were manually set to $\lambda = 50$ and $k = 16000$.

6 PARAMETER VALUE SELECTION

The method presented here has basically 3 parameters: the regularization parameter λ; the granularity control parameter μ and the scale parameter k that controls the robustness of the HR potential. Their setting depends on the particular application and on the nature of the data. For the case of edge–preserving filtering of piecewise almost–constant images corrupted with additive Gaussian noise with standard deviation σ, it is possible to give heuristic rule to perform this tuning. For this kind of images, the field of first differences, $x_{rs} = f_r - f_s$, is a random field with Gaussian probability distribution (with mean equal to zero and standard deviation $\sigma_x = 2\sigma$) corrupted by the probability distribution of the gray scale steps at the edges. The scale parameter in the robust potential $\rho_H(x_{rs}, k)$ was chosen as $k = 1/\sigma_x^2$ (see Refs. [29] [9] for a complete discussion). A robust estimator of σ_x is based on the median of the absolute deviation (MAD) estimator: $\sigma_x \approx 1.4826 MAD(x)$; where $MAD(x) = median(|x_{rs} - median(x)|)$. If only the magnitude is considered (symmetric case), the problem is reduced to $MAD(x) = median(|x_{rs}|)$; in such case, we can use the L_1- modified norm to compute the MAD estimator:

$$\sigma_x \approx 1.4826 \left(\arg\min_m \sum_{rs} \rho_h(|x_{rs}| - m; k_2) \right), \tag{21}$$

where $\rho_h(x; k_2)$ is Huber's potential defined in (13); this approximation is valid for small values of k_2; we fix this value to $k_2 = 0.1$ of the dynamic range of the image and by setting $\lambda = 30$ (a relatively large value), we enforce the constraint of piecewise almost constant images. We found that, for the regularization potential used, the scale parameter computed with (21) overestimates the edges. We control this effect by setting $\mu = 0.01$, and computing the scale parameter k with the empirical rule:

$$k = k(\mu) = \frac{1 + 50\mu}{2\sigma_x^2}. \tag{22}$$

Figure 7 shows the results of performing an edge preserving filtering on images of different kinds. First row shows the original images; second row shows the filtered images and third row shows the detected edges. The parameters were $\lambda = 30$, $\mu = 0.01$ and $k = k(\mu)$ [see (22)]. For the images in columns (a) to (d), k was equal to: 1037, 1263, 2634 and 489, respectively.

As figure 7 shows, this prescription for setting the parameter values works reasonable well in the general case: for particular tasks, however, one may want to change interactively the appearance of the filtered images. Ideally, one would like to do this by moving a single parameter, or better still by monitoring the evolution of the iterative method and stopping it when one gets the desired solution (as in the case of robust anisotropic diffusion [9]). To do this, we introduce a positive parameter c and set $\lambda' = c\lambda$, $\mu' = c\mu$ and $k = k(\mu')/c$; where the apostrophe denotes the updated parameters. If $c > 1$, then the granularity and the detected edges are reduced with respect the filtered image with the original set of parameters; oppositely, for $0 < c < 1$, the granularity and the detected edges are increased. First row in figure 8 shows a sequence of filtered images computed with c equal to 1, 1.5, 2.0, 2.5,and 3.0; respectively. The second row shows the evolution of the iterative algorithm HQCG for $c = 3$. The iteration number for the images in the temporal sequence are 15, 42, 115, 210 and 550, respectively. Note the similarity between the image sequence generated with different values of the hyper–parameter c and the one that depends on time. This allows one to observe this evolution and stop by hand when a solution that is perceptually "optimal" is reached even when the algorithm has not converged, with the significant computational time reduction. One can also note in this experiment that the algorithm presents the nice property of being "continuous" with respect to the parameters, in the sense that a small variation in the parameters produces a small variation in the result.

7 CONCLUSIONS

The main contributions in this paper are the following:

1. We introduced a new class of "Quasi–robust" potentials (QR) for edge–preserving regularization, that are formed by a convex combination of a robust [hard–redescending (HR)] potential and a quadratic one. We have shown that this allows one to control both the edge preservation and the granularity of the solution in a more accurate way. This parametric family of potentials performs a smooth transition from the non–robust (convex) to the HR case as the μ parameter moves from 1 to 0. This allowed us to design a novel continuation parameter method (a form of " Graduate Non-Convexity" [15]) which is particularly effective for the computation of piecewise smooth optical flow.

2. We presented a novel non–linear minimization method (the HQCG algorithm) for cost functions based on half–quadratic potentials. The HQCG is a modification of the non–linear conjugate gradient (NLCG) method and uses a linear CG step to compute the optimal step size in the NLCG procedure. This permits one to obtain restorations that are better of quality that those obtained with other algorithms considered as the state of the art at a fraction of the computational cost.

3. We presented a heuristic method for the selection of parameter values for an important class of problems

(edge preserving filtering of piecewise smooth images).

4. We presented a practical procedure, that operates in a similar way as anisotropic diffusion [9], for interactive edge preserving filtering.

In the experiments presented here, we have only tested the effect of the QR potential in the regularization term. However, it is also possible to use a hard–redescending robust potential ρ_w in the data term (as in [7]) to further improve the solution, particularly for high noise levels.

The added quadratic term in the regularization potential acts as a stabilizer of the conjugate gradient algorithm avoiding the ill–conditioning of the resulting non–linear system because the weight functions (for the data and regularization terms) cannot be both zero (or very small) at the same time. The combination of the QR potentials and the HQCG algorithm allows one to have a more graceful evolution of the solution at each iteration.

Acknowledgement

The authors were partially supported by CONACYT (Mexico) under grant 34575–A.

References

[1] S.Z. Li, *"Markov Random Field Modeling in Image Analysis,"* Springer Verlag, Tokyo, 2001.

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

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

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

[5] A. Rangarajan and R. Chellappa, "A continuation method for image restoration using the adiabatic approximation," in *Markov random Fields: Theory and Applications*, R. Chellappa and A.K. Jain (eds.) Academic Press Inc., 69–91, 1993.

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

[7] M.J. Black and A. Rangarajan, "Unification of line process, outlier rejection, and robust statistics with application in early vision," Int. J. Comput. Vis., vol. 19, no. 1, 57-92 1996.

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

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

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

[11] M. Rivera and J.L. Marroquin, "Adaptive rest condition potentials : First and second order edge–preserving regularization," To Appear in Computuer Vision Image Understanding

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

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

[14] G. Li "Robust Regression," in *Exploring Data Tables, Trends, and Shapes,* D.C. Hoaglin, F Mosteller, J.W. Tukey (eds.), 281–343, Wiley, New York, 1985.

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

[16] J.E. Dennis and R.E. Welsch, "Techniques for nonlinear least squares and robust regression," Proc. Amer. Statis. Ass., 83–87, 1976.

[17] F.R. Hampel E.M. Ronchetti P.J. Rousseeuw and W.A. Stahel, "Robust Statistics: The approach based on Influence Functions," Wiley, New York, 1986.

[18] C.V. Stewart, "Robust parameter estimation in computer vision," SIAM Review, 41, 513–537, 1999.

[19] M.J. Black, D.J. Fleet and Y. Yacoob, "Robusty estimating changes in image appearance," Comput. Vision Image Understand. 78, 8–31, 2000.

[20] M. Ben–Ezra, S. Peleg and M. Werman "Real–time motion analysis with linear programming," Comput. Vision Image Understand. 78, 32–52, 2000.

[21] S.-H. Lai "Robust Image matching under partial occlusion and spatial varying illumination changes," Comput. Vision Image Understand. 78, 84–78, 2000.

[22] P. Hellier, C. Barillot, É. Mémin, P. Pérez, "An energy–based framework for dense 3D registration of volumetric brain images," in Proc. Computer Vision and Pattern Recongnition (CVPR–00), vol 2, 270–275, Hilton Head Is. SC., 2000.

[23] W.J.J. Rey, "*Introduction to Robust and Quasi–Robust Statistical Methods*," Springer–Verlag, Berlin, Heidelberg, 1983.

[24] P.E. Gill, W. Murray and M.H. Wright, "*Practical Optimization*," Academic Press London, New York, 1981.

[25] C.T. Kelley,"Iterative Methods for Linear and Nonlinear Equations," Frontiers in Applied Mathematics, vol 16, SIAM, Philadelphia, 1995.

[26] J.J. Moré and D.J. Thuente, "Line search algorithms with guaranteed sufficient decrease," ACM Trans. Math. Softw. vol 20, 286–307, 1994.

[27] J.C. Gilbert and J. Nocedal,"Global convergence properties of conjugated gradient methods for optimization," SIAM J. Optim. vol 2, 21–42, 1992.

[28] P.K. Horn and B.G. Shunck, "Determining optical flow," Artificielle Intell., vol 17, 185–203, 1981.

[29] P.J. Rousseeuw and A.M. Leory, *Robust Regression and Outlier Detection*, John Wiley & Sons, 1987.

FIGURE CAPTIONS

Figure 1.Potential ρ_Q, influence ψ and weight ω functions with explicit control over the edge definition and granularity.

Figure 2. Segmentation of brain/non–brain tissue.(a) axial section of a MRI. (b) Restoration with HQGN with maximum granularity control ($k = 80$ and $\mu = 0$). (c) Restoration with HQGN algorithm with k is decreased ($k = 60$ and $\mu = 0$). (d) Controlling granularity: restoration with HQGN algorithm with $k = 800$ and $\mu = 0.03$. (e) Boundaries in panel d. (f) Filtered image with the ARTUR algorithm using the Geman–McClude potential.

Figure 3. Vector–valued image processing (channels red, green and blue of a color image). First row, noisy data. Second row, independently computed channels with a HR–potential. Third row, coupled computed channels with a HR–potential. Fourth row, coupled computed channels with a HR–potential and granularity control (see text).

Figure 4. (a) Real image. (b) Filtered with HQCG after 30 iterations (6 secs.). and (c) after 70 iterations (14 secs.). (d) Filtered with ARTUR after 300 iterations (40 secs.).

Figure 5. Performance of the algorithms HQCG (bold line) and ARTUR–CG. Values of the cost function versus iteration number.

Figure 6. Optical flow. (a) and (b) Synthetic image pair (see text for details). (c) and (d) Magnitude of the OF obtained with the quadratic potential and (d) the proposed annealing procedure, respectively.

Figure 7. Test of the method for automatic parameter selection in the filtering task: first column shows the original images, second row the filtered images and third row the detected edges. The computed k value and size in pixels of the images correspond to: (a) 1037 (181×217), (b) 1263 (256×256), (c) 2634 (256×256) and (d) 489 (448×448). In all the cases we used $\lambda = 30$ and $\mu = 0.01$.

Figure 8. First row, images filtered by varying the hyper–parameter c, see text. Second row, image sequence that shows the evolution of the result for $c = 3$.

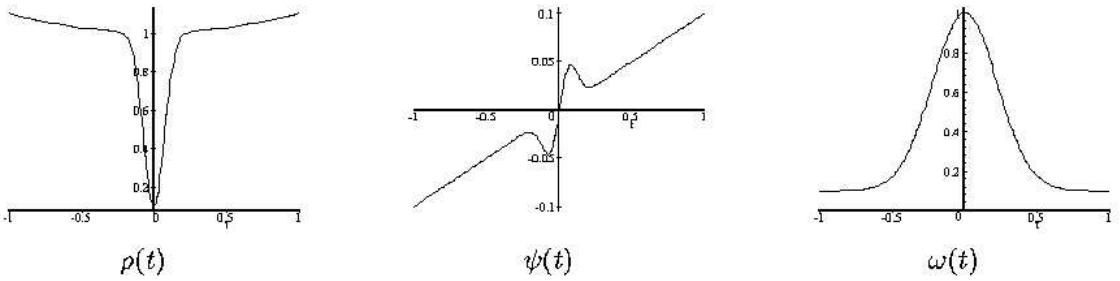

Figure 1:

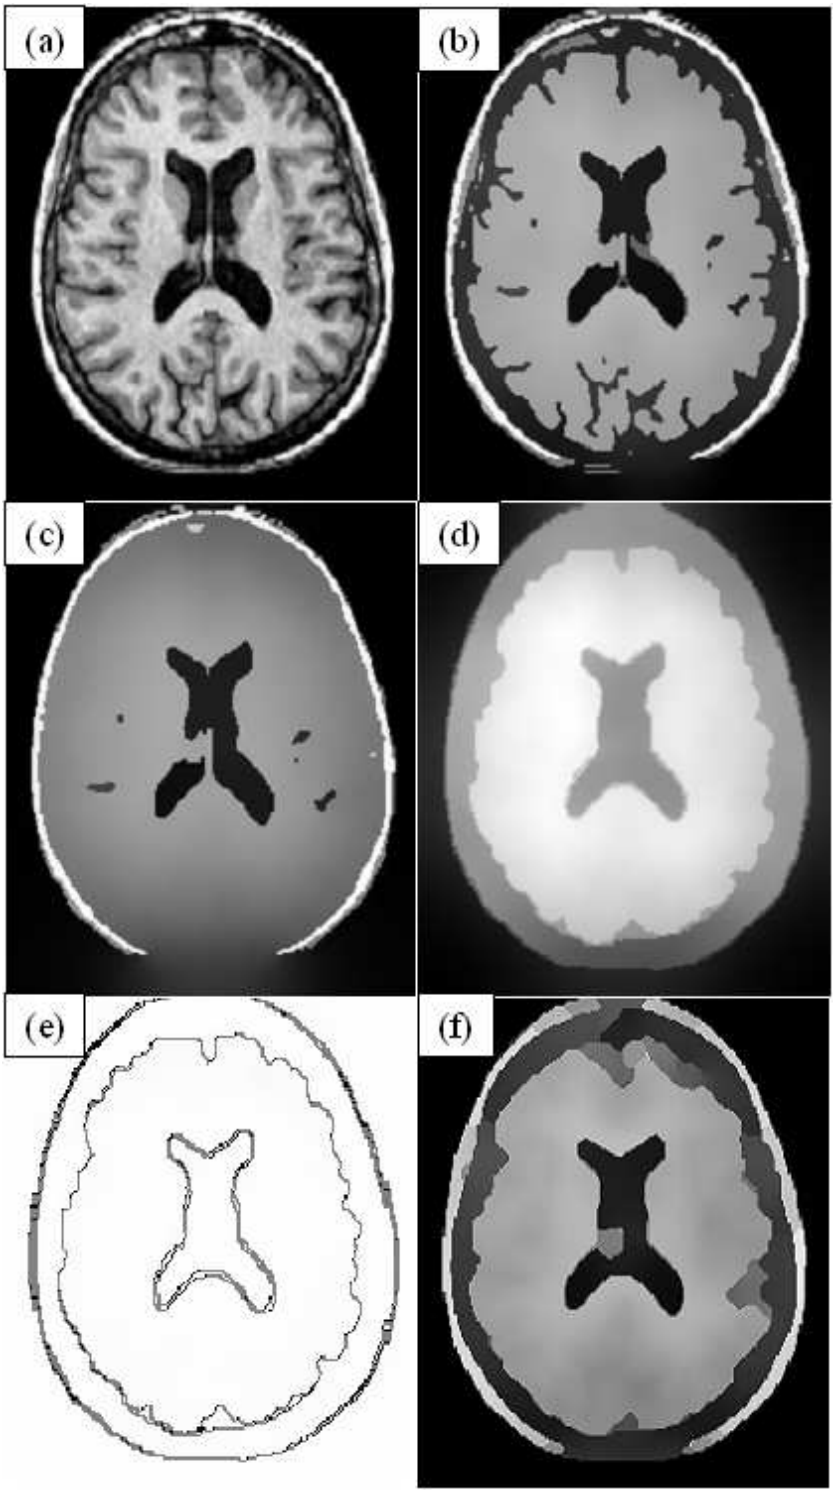

Figure 2:

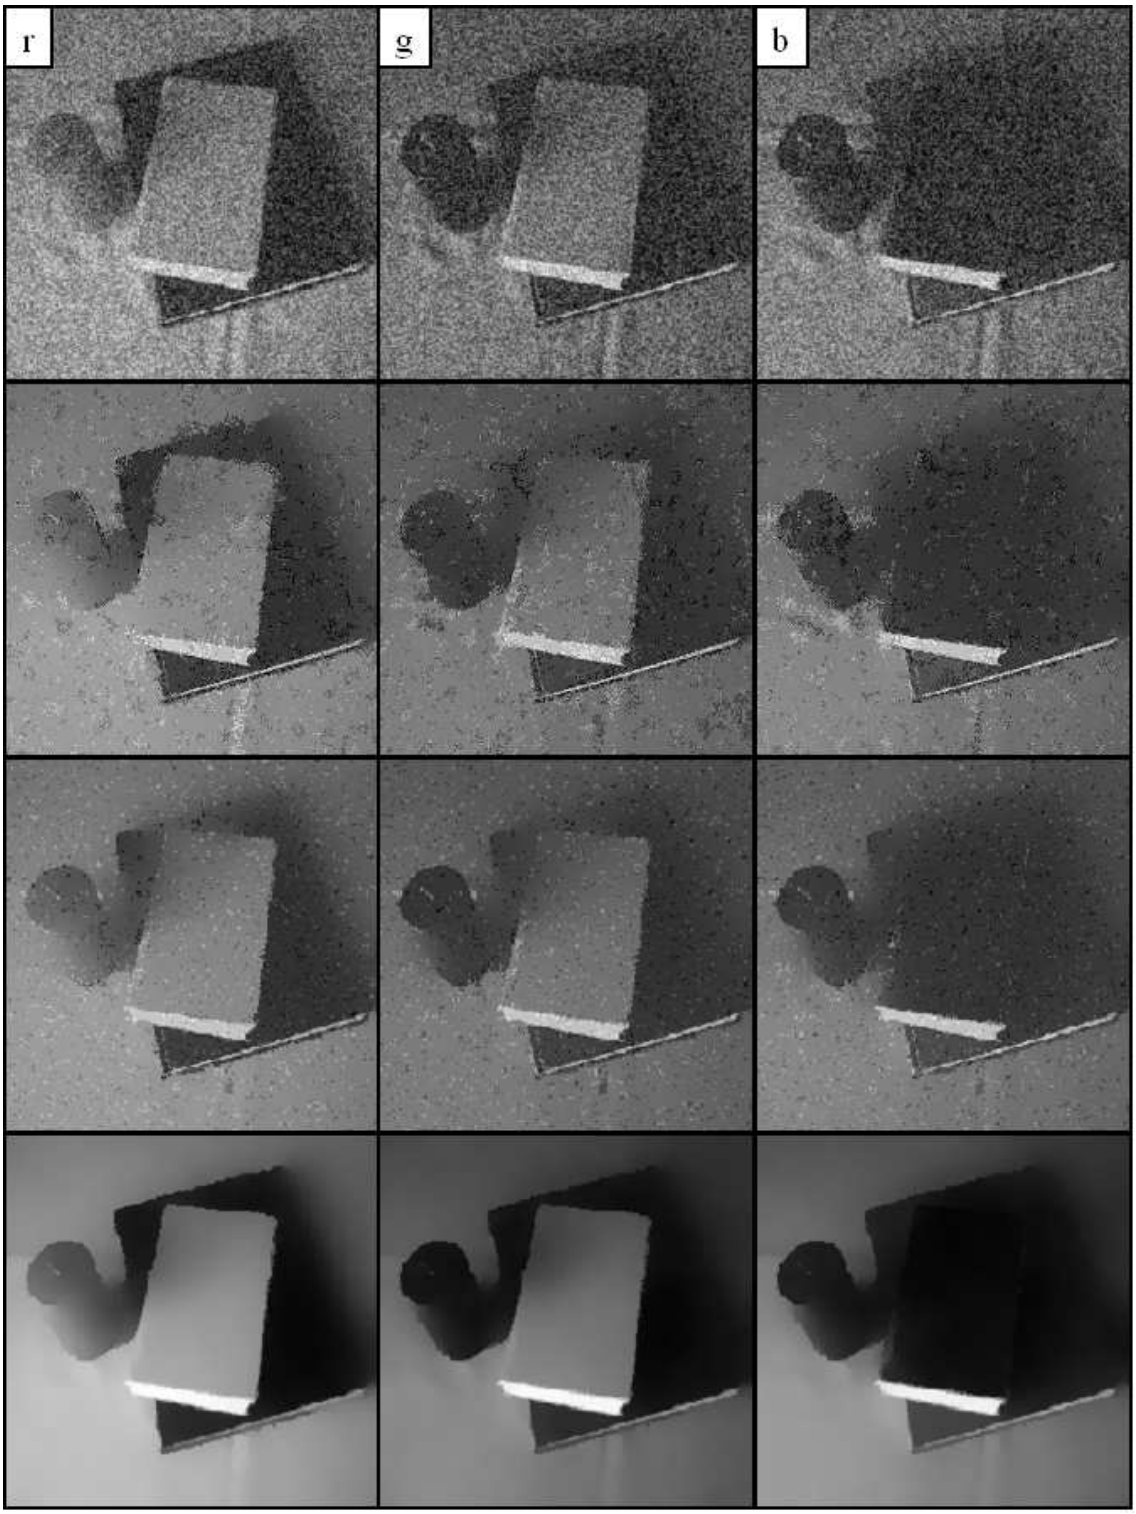

Figure 3:

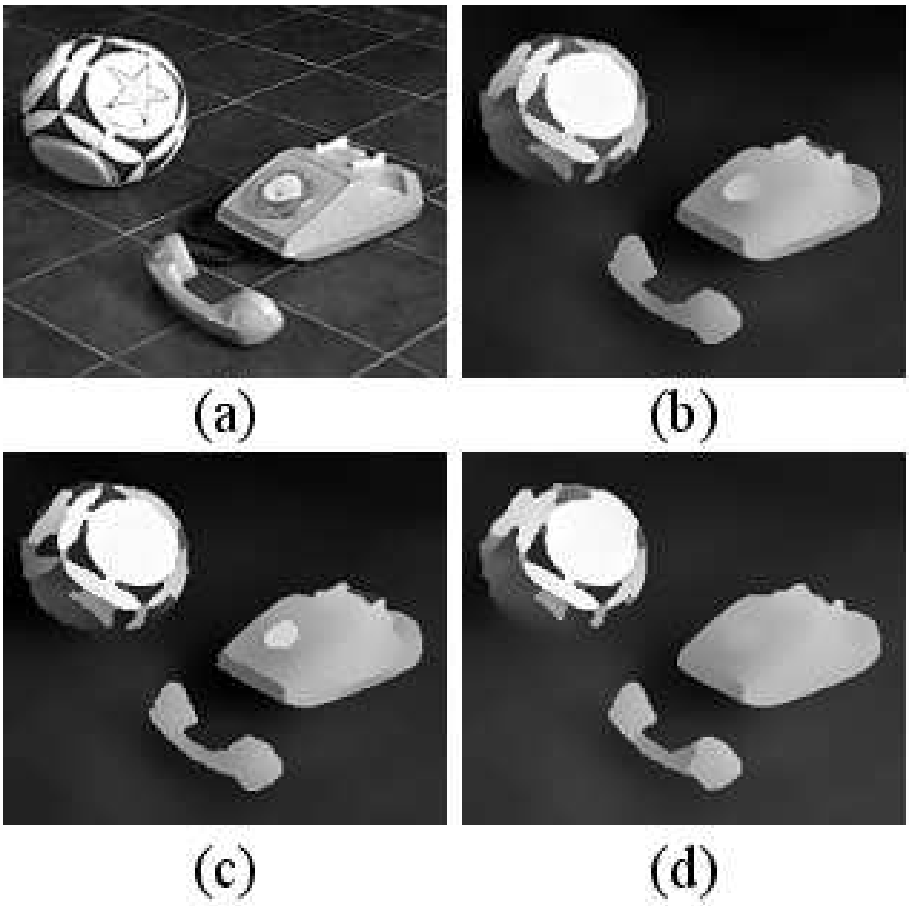

Figure 4:

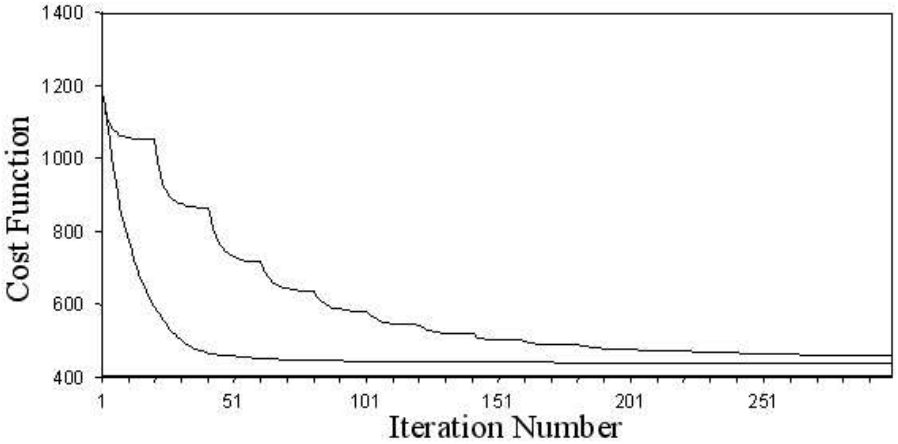

Figure 5:

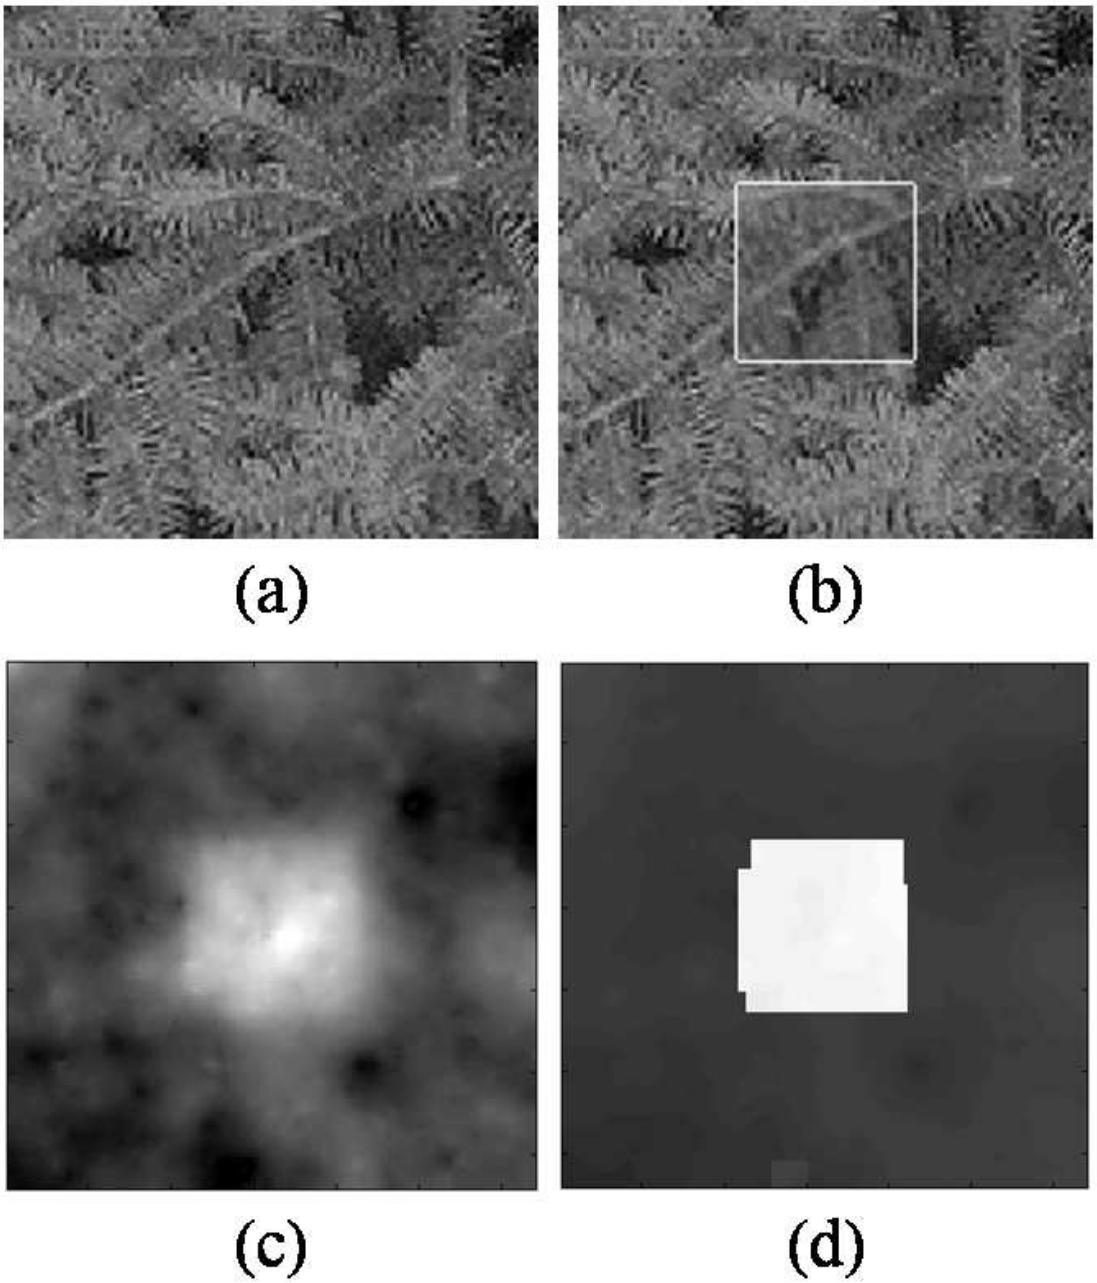

Figure 6:

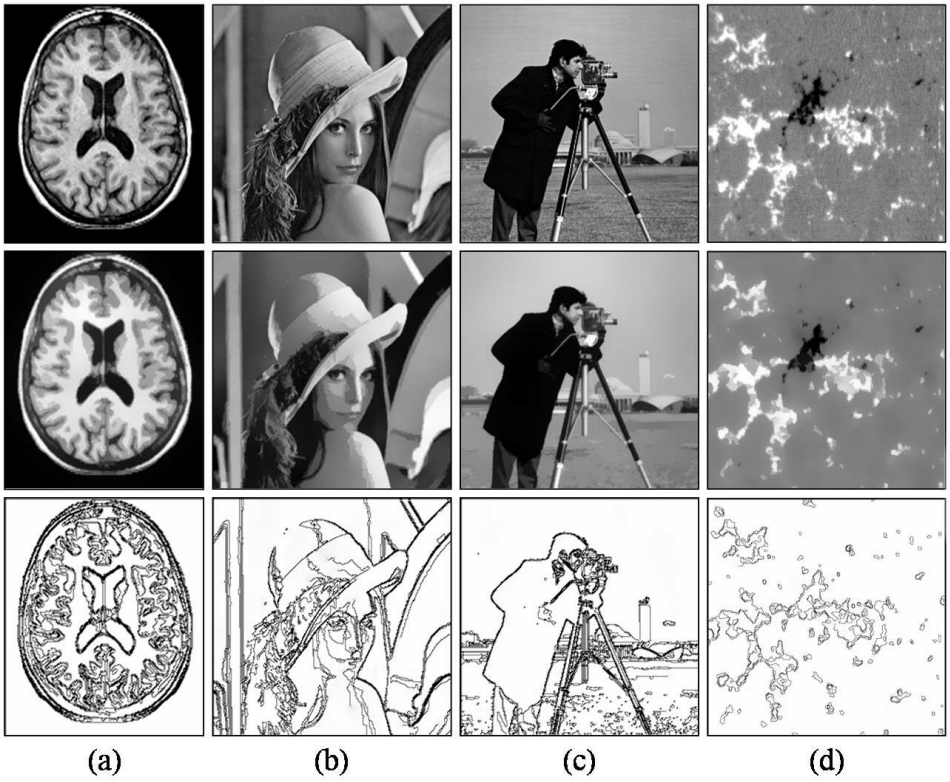

Figure 7:

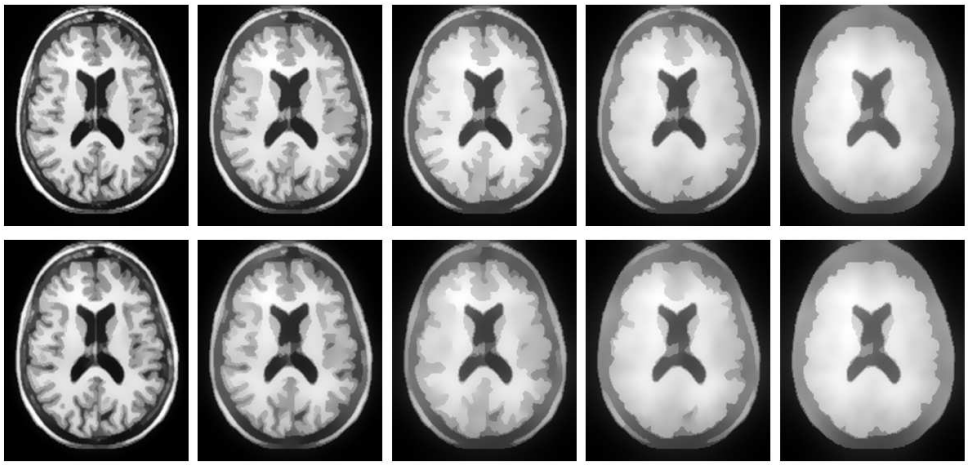

Figure 8: