NIH Public Access

Author Manuscript

Med Image Anal. Author manuscript; available in PMC 2009 April 1.

Published in final edited form as:

Med Image Anal. 2008 April ; 12(2): 87–98.

Image Filtering via Generalized Scale

Andre Souza[1], **Jayaram K. Udupa**[1], and **Anant Madabhushi**[2]

1*Medical Image Processing Group, Department of Radiology, University of Pennsylvania*

2*Department of Biomedical Engineering, Rutgers The State University of New Jersey, 617 Bowser Road, Room 102, Piscataway, NJ 08854.*

Abstract

In medical imaging, low signal-to-noise ratio (SNR) and/or contrast-to-noise ratio (CNR) often cause many image processing algorithms to perform poorly. Postacquisition image filtering is an important off-line image processing approach widely employed to enhance the SNR and CNR. A major drawback of many filtering techniques is image degradation by diffusing/blurring edges and/or fine structures. In this paper, we introduce a scale-based filtering method that employs scale-dependent diffusion conductance to perform filtering. This approach utilizes novel object scale information via a concept called generalized scale, which imposes no shape, size, or anisotropic constraints unlike previously published ball scale-based filtering strategies. The object scale allows us to better control the filtering process by constraining smoothing in regions with fine details and in the vicinity of boundaries while permitting effective smoothing in the interior of homogeneous regions. A new quantitative evaluation strategy that captures the SNR to CNR trade-off behavior of filtering methods is presented. The evaluations based on the Brainweb data sets show superior performance of generalized scale-based diffusive filtering over two existing methods, namely, ball scale-based and nonlinear complex diffusion processes. Qualitative experiments based on both phantom and patient magnetic resonance images demonstrate that the generalized scale-based approach leads to better preservation of fine details and edges.

Keywords

Anisotropic diffusion; image filtering; MR imaging; local scale; filtering quality assessment

1. INTRODUCTION

Noise is ubiquitous in acquired images, especially medical. In many tasks, the utility of an image is determined by how well real intensity interfaces and fine details are preserved in the acquired image. Often the details may be along an edge or a boundary. In addition to influencing diagnostic tasks, noise also affects many image processing and analysis tasks such as segmentation (Pal and Pal, 1993; Saha and Udupa, 2001), registration (Wells et al., 1996; Lester and Arridge, 1999), and visual rendering (Höhne et al., 1990; Udupa and Herman, 2000) that are crucial in many applications.

Address for correspondence: Jayaram K. Udupa, Medical Image Processing Group, Department of Radiology, University of Pennsylvania, Fourth Floor, Blockley Hall, 423 Guardian Drive, Philadelphia, PA 19104-6021, Ph: (215) 662-6780, Fax: (215) 898-9145, E-mail: jay@mipg.upenn.edu.

Methods for improving SNR and CNR can be divided into two categories: those based on image acquisition techniques and methods based on postacquisition image processing. Improving image acquisition techniques usually requires increasing the overall acquisition time, losing spatial resolution, or upgrading the scanner. Filtering, an off-line image processing approach, is often as effective as improving image acquisition without affecting spatial resolution, and if properly designed, requires less time, and is usually less expensive. Filtering techniques may be classified into two families: (i) enhancing, wherein wanted (structure) information is enhanced, hopefully without affecting unwanted (noise) information, and (ii) suppressing, wherein unwanted information is suppressed, hopefully without affecting wanted information. We focus in this paper on filtering as in (ii). Suppressive filtering operations may be further divided into two classes: a) space-invariant filtering (Rosenfeld and Kak, 1982), b) space-variant filtering (Lee, 1980).

Space-invariant filtering techniques, wherein a spatially independent fixed smoothing operation is carried out over the entire image, blur important structures along with noise. For overcoming this problem, space-variant filtering, wherein the smoothing operation is modified by local image features, have been proposed including edge-adaptive filtering (Lev et al., 1977), local image-statistics-driven filtering (Lee, 1980), gradient inverse-weighted filtering (Wang et al., 1981; Wang, 1992), gradient-controlled anisotropic diffusive filtering (Perona and Malik, 1990), wavelet-based multiscale filtering (Xu et al., 1994; Bao and Zhang, 2003), local shape-based template-matched adaptive filtering (Ahn, 1999), nonlinear feature-dependent total variation filtering (Rudin et al., 1999; Chan et al., 2001), fractal based filtering (Ghazel et al., 2003), and kernel principal component analysis based filtering (Kim et al., 2005), and coherence-enhancing diffusion filtering (Weickert, 1999; Tschumperle and Deriche, 2005). Among these, diffusive filtering methods (Perona and Malik, 1990) have become quite popular. In these methods, image intensity at a pixel is made to diffuse to neighboring pixels in an iterative manner, with the diffusion conductance controlled by local intensity gradients. Diffusive filtering techniques have been employed in various applications including filtering of 3D and multi-echo MRI (Gerig et al., 1992; Sapiro and Tannenbaum, 1994), SPECT/PET (Beekman et al., 1998; Demirkaya, 2002), 3D ultrasound (Montagnat et al., 2003), CT (Saha and Udupa, 2001; Krisian, 2002), fMRI (Kim, et al., 2005), visualization (Preuer and Rumpf, 1999; Clarenz, 2000; Meijering et al. 2002), and segmentation (Atkins and Mackiewich, 1998; He and Narayana, 2002).

Perona and Malik’s seminal work paved the way for a variety of diffusive filtering methods that have attempted to overcome some drawbacks and limitations of the original model (see (ter Haar Romeny, 1994) and the references therein). (Catte et. al., 1992) proposed a simple modified diffusion conductance controlled by intensity gradients estimated from a smoothed image. (Weickert, 1998) proposed an anisotropic, nonlinear diffusion process that used both magnitude and direction of gradients. (Liang and Wang, 1998) proposed a piecewise linear diffusion process controlled by a space-variant, anisotropic local noise estimate. (Gilboa et al., 2002) proposed a forward-and-backward adaptive diffusion process that attempted to enhance features while locally suppressing noise in the image.

As we see from the above discussion, the various improvements in diffusion filtering are at better control and adaptation of the diffusive process to local image characteristics. A recent, somewhat different type of attempt towards this goal is through the use of local image scale (Saha and Udupa, 2001; Saha, 2005). Several definitions of local image scale are available in the literature (Elder and Zucker, 1998; Liang and Wang, 1998; Saha et al., 2000; Saha, 2005). In (Saha and Udupa, 2001), scale at any image element is considered to be the radius of the largest ball centered at that element such that all elements within the ball satisfy a pre-defined homogeneity criterion. We will refer to this scale model as *ball scale* (or *b-scale* for short) from now on. The idea behind b-scale-based filtering is to take into account the local scale information at every image element to adaptively control diffusion and the extent of filtering (see Figure 1). The work presented in the present paper is closely related to this idea, in the sense that it uses the spirit of local scale to control diffusion, but it uses a different local/global scale to arrive at improved filtering.

Our proposed filtering method is based on a novel scale idea called, *generalized scale* (*g-scale* for short) (Madabhushi et al., 2006). The g-scale at any image element *p* is considered to be the set of all image elements within the largest, homogeneous, fuzzily connected region containing *p*. Roughly speaking, the g-scale at *p* is the largest set (of any shape whatsoever) of elements within which there is a spatial contiguity of intensity homogeneity. g-scale differs from b-scale and other local morphometric scale models in that it imposes no shape, size, or anisotropic constraints on the homogeneous regions. Filtering is controlled by g-scale regions so that the diffusion rate is greater deep within the region than in its border (see Figure 1)

The nonlinear diffusion filtering process originally proposed by Perona and Malik does not offer any image-dependent guidance for selecting the optimum gradient magnitude at which the diffusion flow must have a maximum value. More importantly, since it does not use any morphological or structural information to control the extent of diffusion in different regions, fine structures often disappear and fuzzy boundaries are further blurred upon filtering. To overcome these problems, b-scale-based filtering was proposed to adaptively control the degree of smoothing that is done in different regions of the image. By definition, the b-scale in regions with fine details or in the vicinity of boundaries is small. Thus, a restricted parameter is automatically selected for filtering in small b-scale regions (corresponding to fine details and the vicinity of boundaries) and a generous filtering parameter is employed at large b-scale regions (corresponding to interiors of large homogeneous regions). While b-scale-based filtering has demonstrated significant improvements in preserving boundary sharpness and fine details while suppressing noise, this (ball) scale model is not appropriate for effective diffusion *along* edges and *along* elongated structures. That is, it does not take into account orientation (shape) and anisotropy of local structures. The motivation for g-scale-diffusion filtering method is to bring in explicitly the object shape, size, and anisotropy information to better control the diffusion process *along* edges and elongated structures (see Figure 2). The shape of the nonlinearity of the diffusion conductance function varies in this method with the g-scale region. A g-scale region may successfully represent a fuzzily connected edge segment or an elongated structure in the image leading to intense diffusion along this structure but not across it.

The rest of this paper is organized as follows. In Section 2, we briefly give an overview of anisotropic diffusion theory and describe briefly a recently published method based on nonlinear complex diffusion (Gilboa et al., 2004) as well as the b-scale-based method (Saha and Udupa, 2001) against which the new method will be compared. The g-scale-diffusion filtering method is presented in Section 2. In Section 3, results of qualitative and quantitative evaluations on MR and CT clinical and phantom images are presented by comparing among the three methods by using a new evaluation strategy that considers the tradeoff that exists between CNR and SNR. We state our b-scale regions g-scale regions conclusions in Section 4. A preliminary version of this paper was presented at the SPIE Symposium on Medical Imaging 2005 (Souza et al., 2005).

2. METHODS

For brevity, we refer to an acquired digital volume image as a *scene* and represent it by a pair $\mathcal{C} = (C,f)$ where $C=\{c \mid -b_j \le c_j \le b_j \text{ for some } b \in Z^3_+\}$, Z^3_+ is the set of 3-tuples of positive integers called *voxels*, *f* is a function whose domain is *C*, called the *scene domain*, and whose range is a set of integers [*L*, *H*], and for any $c \in C$, $f(c)$ is referred to as the *intensity* of *c*. We call C a *binary scene* if the range of f is {0,1}. A *digital ball* (or simply a *ball*) of radius r centered at any voxel c in C is the set $B_r(c)=\{d \in C| \|c - d\| \leq r\}$. For any set X, we use the notation $|X|$ to denote its cardinality.

In the rest of this section, we will first outline the nonlinear diffusion method (Perona and Malik, 1990), then discuss its extension as employed in the method of nonlinear complex diffusion (Gilboa et al., 2004), and then describe the b-scale-based and g-scale-based methods.

2.1. Anisotropic diffusion

As described in (Perona and Malik, 1990), anisotropic diffusion is a locally adaptive smoothing process, which attempts to minimize blurring near object boundaries. A mathematical formulation of the diffusion process on a vector field **V** at a point c in coordinate-free form can be given by

$$\frac{\partial f}{\partial t}=div\,\mathbf{V}=\lim_{\Delta\tau\to 0}\int_s \mathbf{V}\cdot d\mathbf{s}, \tag{1}$$

where $\Delta\tau$ is the volume enclosed by the surface s (surrounding c) and $d\mathbf{s}=\mathbf{u}ds$, where $\mathbf{u}$ is a unit vector which is orthogonal and outward-directed with respect to the “infinitesimal” surface element ds The intensity flow vector field **V** controls the diffusion process, and is defined as

$$\mathbf{V}=G\mathbf{F}, \tag{2}$$

where G is the diffusion conductance function, and **F** is the scene intensity gradient vector field. In a linear isotropic diffusion process, G is a constant, and in (Perona and Malik, 1990), the authors have argued that such diffusion strategies blur object boundaries and structures. They presented an alternative, anisotropic diffusion method in which G varies at each location in the scene as a nonlinear function of the magnitude of the scene intensity gradient so that smoothing within a region with low intensity gradients is encouraged, and smoothing across boundaries, wherein the magnitude of the gradients is much higher, is discouraged. Although diffusion is a *continuous* process, in image processing, diffusion is achieved by an iterative process in the *discrete* domain. Our method is formulated in 3D discrete spaces by using the 6-adjacent voxel neighborhood. This dichotomy between continuous versus discrete representations of the diffusion process is not new and has already been addressed by earlier papers. The core idea of using scale (either b- or g-scale) is to finely control the amount of diffusion locally by fine tuning the conductance parameter rather than to assume it to be constant for the whole image as proposed by Perona and Malik.

2.2. Nonlinear complex diffusion

The *nonlinear complex diffusion* (NCD) method of (Gilboa et al., 2004) presents a complex diffusion-type process, which generalizes the nonlinear anisotropic diffusion process by incorporating the free Schrödinger equation. An important characteristic of the NCD method is that the imaginary part of the NCD process acts as an edge detector (smoothed second derivative scaled by the phase angle θ and time) when the complex diffusion-conductance function approaches the real axis (i.e., when θ is very small). The real part of the NCD process is effectively decoupled from the imaginary part, and behaves like a real nonlinear diffusion process. The authors have shown that NCD has better performance than the Perona and Malik diffusion process in overcoming staircasing effect and in handling changes in the scene illumination condition. Based on this property, they devised a complex functional form for the diffusion-conductance function G_n, given by

$$G_n(c)=\frac{e^{i\theta}}{1+\left(\frac{Im(f_t(c))}{k\theta}\right)^2}, \tag{3}$$

where for any $c \in C$ $Im(f_t(c))$ is the imaginary part of the complex intensity value $f_t(c)$ resulting from the complex diffusion process of NCD at tth iteration, and k is a threshold parameter for the gradient magnitude. The phase angle θ of the complex function G_n does not change over time (iterations). Since the imaginary part is normalized by θ, the process is almost not affected by changing the value of θ, as long as it stays small (θ < 5°). $Im(f_t(c))$, for any $c \in C$, is assumed to be zero at $t = 0$ (first iteration). A discrete implementation is given in (Gilboa et al., 2004). Diffusion flow magnitude function |**V**| has a maximum value at magnitude of gradient |**F**| = k, |**V**| is monotonically increasing for |**F**| < k, and it is monotonically decreasing for |**F**| > k. When k is large, the generosity of filtering is high and the possibilities of blurring across boundaries increase. On the other hand, for small k, the generosity of filtering is low. Because k is fixed in this method while filtering a given scene as in (Perona and Malik, 1990), a fine control on and adaptivity to local region characteristics are missing. How to bring this local control and adaptivity is the main thrust of this paper.

2.3. b-scale-based diffusive Filtering

For any ball $B_r(c)$ of radius r centered at c, we define a fraction $FO_r(c)$, that indicates the fraction of the set of the voxels in the ball boundary whose intensities are sufficiently uniform with that of c, by

$$FO_r(c)=\frac{\sum_{d\in B_r(c)-B_{r-1}(c)} W_\psi\left(|f(c)-f(d)|\right)}{|B_r(c)-B_{r-1}(c)|}, \tag{4}$$

where $W_\psi(x)$ is a fuzzy membership function corresponding to the predicate "x is small". A complete discussion of the characteristics of W_ψ is presented in (Saha et al., 2000). Basically, W_ψ should satisfy the following properties: function range should be [0,1], $W_\psi(0)=1$, and should be monotonically decreasing with $|f(c) - f(d)|$. In this paper, a zero mean unnormalized Gaussian function with standard deviation σ_ψ is used for W_ψ. σ_ψ is a homogeneity parameter which, for the purposes of the present paper, is automatically estimated from the given scene as described in (Saha and Udupa, 2001). σ_ψ is therefore not considered to be a parameter of the bD filter.

The algorithm for b-scale estimation, described in (Saha et al., 2000), iteratively increases the ball radius r by 1, starting from r=1, at every voxel $c \in C$, and checks $FO_r(c)$, the fraction of the object containing c that is contained in the ball boundary. The first time when this fraction falls below the tolerance parameter t_s, the ball is considered to enter a region of different homogeneity from that to which c belongs. Following the recommendation in (Saha et al., 2000), we have used $t_s = 0.85$ in this paper. The output of the algorithm is a *b-scale scene* $C_S = (C,f_S)$, where $f_S(c)$ is the radius of the largest ball centered at c within which the voxel intensities are homogeneous.

The *b-scale-based diffusion* (bD) method is iterative, and so let t denote the iteration number. Let $C_t = (C,f_t)$ denote the scene resulting from the diffusion process at tth iteration. For any voxels $c,d \in C$, such that $c \neq d$ and c and d are 6-adjacent, let $\mathbf{D}(c,d)$ denote the unit vector along the direction from voxel c toward voxel d. Let $\mathbf{F}_t(c,d)$ represent the component of the intensity gradient vector along, $\mathbf{D}(c,d)$ given by

$$\mathbf{F}_t(c,d)=\frac{f_t(c)-f_t(d)}{\sqrt{\sum_{i=1}^{3}\frac{\nu_i^2(d_i-c_i)^2}{\min_j[\nu_j^2]}}}\mathbf{D}(c,d). \tag{5}$$

Note that scene resolution anisotropy is taken into account in (5), where $\nu = (\nu_1,\nu_2,\nu_3)$ is the voxel size.

The intensity flow from voxel c to voxel d is affected by the *effective scale* $r_{eff}(c,d)$, which is defined as

$$r_{\text{eff}}(c,d)=\min[\,f_s(c),f_s(d),f_s(e)], \tag{6}$$

where $e \in C$ is the neighboring voxel of c just opposite d, defined by $e = c - (d - c)$. (When c is in the border of the scene domain, we let $e = c$.) The diffusion conductance function for the flow from c to d at the tth iteration, via this definition of effective scale, is defined as follows:

$$G_t(c,d)=e^{-\frac{|F_t(c,d)|^2}{2\{\sigma_s(c,d)\}^2}}, \tag{7}$$

where $\sigma_s(c,d)$ is the b-scale-adaptive region-homogeneity parameter used for controlling the intensity flow from d to c, given by

$$\sigma_s(c,d)=\frac{\sigma_\psi\{1+r_{\text{eff}}(c,d)\}}{1+r_{\text{MAX}}}. \tag{8}$$

$G_t(c,d)$ is an unnormalized Gaussian and a valid conductance function as proposed by (Perona and Malik, 1990). The intention here is to let $G_t(c,d)$ tend to zero at the edges where b-scale (and hence $r_{eff}(c,d)$) is small and so $G_t(c,d)$ decreases. In all the experiments in this paper, the maximum ball radius r_{MAX} in the b-scale scene C_S of C is set to 12. σ_ψ is the homogeneity parameter defined earlier. Intensity flow vector $\mathbf{V}_t(c,d)$ from voxel c to voxel d at the tth iteration is now defined by

$$\mathbf{V}_t(c,d)=G_t(c,d)\mathbf{F}_t(c,d). \tag{9}$$

Then the iterative process is defined as follows:

$$f_t(c)=\begin{cases} f(c), & \text{for } t=0\\ f_{t-1}(c)-K_D\sum_{d\in C}\mathbf{V}_{t-1}(c,d)\cdot\mathbf{D}(c,d), & t>0, \end{cases} \tag{10}$$

where K_D is the integration constant that must be adjusted according to the adjacency criterion. We use $K_D = 1/7$, as indicated in (Saha and Udupa, 2001) and (Gerig et al., 1992) for 6-adjacent neighborhood. The flow direction between any voxels $c,d \in C$ is always such that it tries to reduce the gradient between them, i.e., $\mathbf{V}_t(c,d)\cdot\mathbf{D}(c,d)$ is positive when $f_t(c) > f_t(d)$, negative otherwise, and zero when $c = d$. Further, this diffusion process described by (5)–(10) is both nonlinear and anisotropic.

2.4. g-scale-based diffusion

As defined in (Madabhushi et al., 2006), for any number $t_g \in [0,1]$, and a homogeneity function W_ψ (as introduced in Section 2.3), the g-scale $G(c)$ of any $c \in C$ in a given scene $\mathcal{C} = (C,f)$, is the *largest* subset of C such that

1. $G(c)$ contains c, and
2. for any α_b-adjacency of voxels in C, say 6-adjacency, and for any voxel $d,o \in G(c)$, there exists an α_b-path $p_{do} = < c^{(1)} = d, c^{(2)}, \ldots, c^{(m)} = o <$ in $G(c)$ connecting d to o such that, for $1 \le I < m$, $W_\psi(|f(c^{(i)}) - f(c^{(i+1)})|) \ge t_g$.

The *g-scale set* $\mathcal{G}(C)$ of C is the family of sets. $\{G(c) \mid c \in C\}$. The g-scale regions collectively cover the entire scene domain but also they do not overlap – that is, any two g-scale regions are either the same region in C or they are completely disjoint. This implies that, for any voxel $d \in G(c)$, there is no need to compute. $G(d)$. The g-scale set of any given scene can be computed by using the following algorithm. As described in (Madabhushi et al., 2006), the two parameters σ_ψ and t_g associated with estimating $\mathcal{G}(C)$ are automatically determined.

Algorithm gSE

Input: C, W_ψ, t_g

Output: $\mathscr{G}(C)$.

Auxiliary Data Structures: A queue Q.

begin

unmark all voxels of C;

repeat

take an unmarked voxel c of C and put it in Q;

mark c and set $G(c)$ to empty;

while Q is not empty *do*

remove a voxel d from Q and add it to $G(c)$;

for each unmarked voxel e in C α_b-adjacent to d *do*

if $W_\psi(|f(d) - f(e)|) \geq t_g$ *then* put e in Q and mark it;

endfor;

endwhile;

output $G(c)$;

until there is no unmarked voxel in C;

end

In *g-scale-based diffusion* (gD), the adaptive region-homogeneity parameter $\sigma_s(c,d)$ is redefined by using g-scale morphometric information as follows.

$$\sigma_s(c,d)=\begin{cases} \frac{\sigma_\psi}{G_{\max}}, & if\ G(c) \neq G(d) \text{ and } |G(c)|>\beta \text{ and } |G(d)|>\beta, \\ \sigma_\psi, & \text{otherwise}, \end{cases} \tag{11}$$

where $G_{\max}$ is the maximum g-scale size $|G(c)|$, and β is a threshold on the size of g-scale. Now the process of g-scale-based diffusion is completely described by (5), (7) and (9)–(11). Intuitively, β should be small because it characterizes potential noise regions in the image, which appear as small g-scale regions, where diffusion should not be constrained (i.e., then $\sigma_s = \sigma_\psi$). This equation assumes that diffusion *across* the boundary of large g-scale regions (> β) should not occur (i.e., then $\sigma_s = \sigma_\psi/G_{\max}$). And $\sigma_s(c,d)$ is equal to σ_ψ inside these regions. We note that, in this process, diffusion takes place along boundaries in a uniform manner, no matter what their shape is, but not across them. For large, β, gD operates as a classical anisotropic diffusion filter. We did not see any significant change in gD results for β values in the range from 5 to 8. Filter performance degraded considerably for β values less than 5 and greater than 8 in our tested images. In both bD and gD, the scale values are evaluated only once at the beginning and not modified in the iterative diffusive process.

3. EXPERIMENTS, RESULTS, DISCUSSION

In this section, we present the results of qualitative and quantitative experimental evaluation by comparing the three methods bD, NCD, and gD. In all experiments, the following values for the parameters of the three methods have been used. For the NCD method, we have used $k = \sigma_\psi$ and $\theta = \pi/60$ in (3) as recommended by the original method. In both gD and bD methods, we have set the tolerance parameter $t_s = t_g = 0.85$ as suggested and justified in the original methods. For the gD method, we have used β = 8 in (11) as explained in the previous section. The data sets utilized in our evaluation are described briefly in Table 1.

3.1 Qualitative Evaluation

Data sets D1–D6 are used for qualitative evaluation. In Figure 3, we display an example from each of these data sets D1–D5 and the filtered results produced by the three methods after three iterations. Filtering has been done in 3D space in all these cases although the results are shown for only one slice. To minimize the number of displays, somewhat zoomed in versions only are shown. Overall, we may observe that the bD and gD methods preserve fine details and edges in the diffusion process better than the NCD method. Note and compare particularly the regions indicated by arrows. At the same time, the interior of homogeneous regions are smoothed well by bD and gD as well as by NCD methods. This effectiveness of bD and gD seems to be independent of the body region and the imaging modality and protocol.

A particular strength of gD over bD (and NCD) is its ability to diffuse well *along* edges, at the same time minimizing diffusion *across* them. Most diffusion based methods, if run for a sufficiently large number of iterations, will eventually not only smooth the interior homogenous regions and along edges but will also smooth across edges. The relative merits of different methods come from their ability to smooth more within object regions and along edges for a given amount of blurring (diffusion) they commit across edges. That is, between two methods M1 and M2, if M1 achieves more diffusion within homogeneous regions and along edges than M2 for the same given amount of blurring committed by both M1 and M2 across edges, then M1 is to be regarded as a better method than M2. In our view, this is what *all* diffusion based (for that matter, all noise suppression) filtering methods try to accomplish. In the next section, we present a new method for evaluating filtering techniques that can properly address this dichotomy. The purpose behind this discussion is to make the point that, in the above sense, gD is a better diffuser than bD (and NCD). Why this is so is intuitively easy to explain. The b-scale regions are very small (1–2 voxel diameter) in the vicinity of object boundaries (see Figure 2). Every voxel along object boundaries is thus likely to be in a different b-scale region, whereas, since g-scale regions have no restriction on shape, a single g-scale region may typically run along a boundary consisting of a run of voxels on the boundary. Therefore, by (11), we can see that gD achieves more diffusion along the boundary also. This is borne out in Figure 3 for data sets D1–D4. D5 is specially created to demonstrate this phenomenon. This phantom was generated by first blurring by applying a 2D Gaussian kernel (σ_{blurr} = 2.0) the original spiral pattern image and subsequently by adding a correlated zero mean Gaussian noise (σ_{noise} = 20). Better smoothing along the edges by gD without introducing blur across edges is evident in Figure 3 for D5 compared to other methods.

PDE based anisotropic diffusion filtering frameworks have been proposed recently that use so called *diffusion tensors* to smooth along edges (Weickert, 1999; Tschumperle and Deriche, 2005). We have tested the Weickert diffusion method on the image data sets used in this paper. Some examples are shown in Fig.4. The Weickert diffusion process via the second-moment matrix has been proved to be useful for images with flow-like texture (e.g., Diffusion Tensor MRI). We find that, after applying Weickert diffusion, the gyri and sulci patterns in brain MR images are distorted and the bone trabecular structures are destroyed, as indicated by arrows in Figure 4. Strictly speaking, the PDE equation by Weickert et al. does not fully respect the tensor geometry. Yet there is a difficult trade-off between noise removal and preservation of curved structures, when using trace-based PDE's (Tschumperle and Deriche, 2005).

To further demonstrate the phenomenon of better smoothing along edges, we present in Figure 5 3D renditions (Udupa and Odhner, 1993) from data set D6. Zero mean correlated Gaussian noise (σ_{noise} = 100) was added to the original CT scene to test the ability of the three methods to suppress noise. The skull surfaces were extracted from the noisy scene and from the three filtered scenes by applying the same fixed threshold. The top row of Figure 5 displays renditions of these surfaces. In the bottom row of Figure 5, we display zoomed in renditions of the teeth and mandible. Again, the sites of interest where important differences exist in filtering effects are indicated by arrows. Surface rendition from scale-based filtered scenes appears notably sharper than the one obtained by NCD method. Yet surfaces appear smoother with gD than with bD filtering. In addition, surface interfaces and fine details of the teeth and mandible are better retained (or even enhanced) by gD than by bD while suppressing noise. The differences between bD and gD methods are difficult to assess and portray visually because of the trade-off issue referred to earlier. This issue has not been addressed in the literature. Published methods show usually at one parameter setting some qualitative/quantitative improvement. This does not convey the full story of the trade-off behavior. Methods should be compared, we argue, based on their entire range of behavior, as described in the next section.

3.2 Quantitative Evaluation

For a quantitative comparison among the methods, we used data set D7 consisting of 45 MRI simulated phantom scenes corresponding to (i) three levels of noise (3%, 7%, and 9%), (ii) three protocols (PD, T1, and T2), and (iii) five slice thicknesses (1 mm, 3 mm, 5 mm, 7 mm, and 9 mm). Let $\{C_{xi} \mid C_{xi} = (C, f_{xi}), 1 \leq i \leq 45\}$ denote the set of scenes produced after applying the filtering method $x \in$ {bD, NCD, gD} to the set of 45 phantom scenes. For a given object region in these scenes and for a given method x, we define *residual noise* RN_{xi} in the scene C_{xi} as the standard deviation of voxel intensities $f_{xi}(c)$ in C_{xi} within the object regions in C_{xi}. (The object that was considered for these data sets was the white matter region. The true tissue region information available in Brainweb as binary scenes was used for this purpose.) Similarly *relative contrast* RC_{xi} of the object regions in C_{xi} is defined as:

$$RC_{xi} = \frac{|M_{xi} - M_{xi}^{\mathrm{IE}}|}{\sqrt{\sigma_{xi}\,\sigma_{xi}^{\mathrm{IE}}}}, \tag{12}$$

where M_{xi} and σ_{xi} denote the mean and standard deviation, respectively, of voxel intensities $f_{xi}(c)$ in C_{xi} within the object region, and M_{xi}^{IE} and $\sigma_{xi}^{\mathrm{IE}}$ denote similar entities outside the object region in their immediate exterior(i.e., a set of all voxels in the background which are 6-adjacent to some voxel in the object). In Figure 6, we demonstrate the behavior of each method by plotting how the *RN* and *RC* values vary as the number of iterations is varied for the different protocols at different levels of noise. The *RN* and *RC* values and the curves shown here are averaged over the five scenes corresponding to five different slice thickness values for each protocol and noise level. The values of RN_{xi} and RC_{xi} were normalized by their maximum observed in these experiments. Table 2 lists the *area under the curve* (AUC) of these values for each curve in Figure 6.

We note that, in assessing the performance of a suppressing filter, the trade-off that exists between *RN* and *RC* should be analyzed. This is akin to the trade-off issue between false positives and false negatives arising in image segmentation and object detection tasks. In analogy with ROC curves for the latter task, we call the curves depicting the trade-off behavior between *RN* and *RC filter operating characteristic* (FOC) curves. As discussed previously, it is usually easy to reduce *RN* simply by increasing the number of iterations or by changing the parameter that controls the conductance function. However, it is challenging to reduce *RN* and at the same time increase/maintain *RC*. The FOC curve depicts how well this trade-off is handled by a filtering method. The upper left corner in this graph represents the ideal filter operating point of least residual noise and highest contrast that is possible. The upper right corner indicates the operating point after a large number (essentially infinite) of iterations when the residual noise reduces to zero but so does the contrast. The AUC value is therefore a good measure of the overall performance of a filter. A higher value of AUC for the method indicates more effective filtering. From Table 2, we observe that, for every protocol and each level of noise, AUC for scale-based diffusive filtering (bD and gD) is higher than that for the NCD method and that the gD method always outperforms the bD method. Continuing our earlier discussion on filtering within uniform regions and along boundaries versus blurring boundaries, we note from Figure 6 that, for every level of noise and for each protocol, gD achieves more smoothing (lower *RN*) for the same level of boundary blur (*RC*) than bD and NCD. Similarly, it achieves less boundary blur (higher *RC*) than bD and NCD for the same level of noise suppression (*RN*). Although, in this paper, we varied only the number of iterations for FOC curves, other important parameters can also be considered for this purpose. In this case, each point on the FOC curve would denote a particular setting of the parameter vector for the method.

One may argue that σ_ψ may not be the optimal value for k in the complex functional form for diffusion-conductance in (3). In Table 3, we list AUC values for all three methods estimated over five different values of σ_ψ given by $0.25\sigma_\psi$, $0.50\sigma_\psi$, $0.75\sigma_\psi$, $0.75\sigma_\psi$, and $1.50\sigma_\psi$, on a MNI phantom scene with 9% noise and 9 mm slice thickness. We observe that variation in NCD filtering performance in terms of change in AUC values was not significant. Both scale-based methods outperformed NCD for all different settings of σ_ψ, while gD had the highest AUC values among the three.

On a Pentium IV (3.4 GHz and 1 GB RAM) PC, the gD method performs as quickly as the NCD method, on average, under 1 minute for three iterations for a 256×256×51 scene. The bD method requires more iterations to achieve a level of filtering similar to that achieved by gD in homogeneous regions. For roughly the same level of filtering in homogeneous regions, bD takes about 2 minutes for the same scene.

4. CONCLUDING REMARKS

In this paper, we have presented a new scale-based diffusive filtering method, called gD method, wherein a fine control on and adaptivity to local region characteristics is incorporated by using a local morphometric scale model, called g-scale. Unlike other local morphometric scale models, g-scale imposes no shape, size or anisotropic constraints. The mechanism underlying the scale-based methods bD and gD is to use the size of the scale region to control diffusion. Equation (8) for bD uses the fact that small values of b-scale (the radius of the largest homogenous ball centered at each voxel) occur in the vicinity of the boundaries where diffusion should be constrained, but allow strong diffusion to take place where b-scale value is large (at the center of homogenous regions). Equation (11) for gD has a clear theoretical advantage as compared with (8), namely that g-scale can have any shape and can potentially drive diffusion *along* boundaries (but not across them) in a highly non-linear manner. This equation assumes that diffusion across different g-scale regions should not occur but should be at maximum inside these regions. Most parameters are computed automatically and the only free parameter in (11), β, is very intuitive. It characterizes the potential noise regions in the image, which appear as small g-scale regions. β indicates what is the minimum region size of g-scale so that for any g-scale region with region size less than β diffusion across boundary is allowed.

We demonstrated that the gD method outperforms both bD and NCD methods and a recent diffusion tensor based method. The filtered scenes produced by the gD method have fine details and structures better preserved while suppressing noise in large homogeneous regions and along boundaries.

The gD method achieves a considerable gain in terms of AUC values over the recently reported NCD method and improves over bD. Computationally it is comparable to NCD and is faster than bD. Thus, gD becomes a powerful contender for filtering noise in medical images at the present time. White matter regions were chosen for evaluation because it occupies a large portion in the brain. While the same analysis can be done for other regions, we have limited ourselves to that region only during evaluation for the following reason. Grey matter and peripheral cerebrospinal fluid are very thin objects so that, RN, which defines the standard deviation of voxel intensities within the object, is difficult to estimate because, in these objects, most voxels are boundary voxels. There is not a sufficient number of voxels in the proper object interior to provide sufficient statistics.

Another contribution of this paper is a methodology to quantify the trade-off behavior between residual noise and residual contrast that exists in all filtering methods. The AUC mechanism proposed here allows filtering methods to be compared over their entire range of behavior rather than at some arbitrary operating point set in an ad hoc manner, as done in most published filtering methods.

While the AUC value is useful in comparing methods over their range of behavior, it does not suggest how to set the parameters of a particular method for a desired operating point. We suggest that by studying the dynamics of the filtering method via FOC, one can devise techniques even for the optimal selection of filter parameters including the number of iterations. The idea consists of graphing the FOC curve in such a way that each operating point on the curve represents one possible choice of the filter parameter vector values including the number of iterations. Our aim is to find that operating point on the curve that represents the best performance. One possible choice of this point may be the point on the curve closest to the upper left corner. In this case, the optimization criterion is the distance of the FOC curve from the upper left corner of the graph. We may minimize this distance as a function of the parameters of the filter to arrive at an optimal operating point for the filter. This approach will fill the void that currently exists in the choice of filter parameters.

The performance of both bD and gD may improve further if scale estimation is done again after each iteration of filtering and if scale estimation and filtering are carried out iteratively.

ACKNOWLEDGEMENT

The research reported here is supported by NIH grants NS37172 and EB004395.

REFERENCES

Ahn CB, Song YC, Park DJ. Adaptive template for signal to noise ratio enhancement in magnetic resonance imaging. IEEE Trans. Med. Imaging 1999;18:549–556. [PubMed: 10463132]

Atkins MS, Mackiewich BT. Fully automatic segmentation of the brain in MRI. IEEE Trans. Med. Imaging 1998;17:98–107. [PubMed: 9617911]

Bao P, Zhang L. Noise reduction for magnetic resonance images via adaptive multiscale products thresholding. IEEE Trans. Med. Imaging 2003;22:1089–1099. [PubMed: 12956264]

Beekman FJ, Slijpen ETP, Niessen WJ. Selection of task-dependent diffusion filters for the post-processing of SPECT images. Phys. Med. Biol 1998;43:1713–1730. [PubMed: 9651035]

Catte F, Lions PL, Morel JM, Coll T. Image selective smoothing and edge detection by nonlinear diffusion. SIAM J. Num. Anal 1992;29:182–193.

Chan TF, Osher S, Shen J. The digital TV filter and nonlinear denoising. IEEE Trans. Image Process 2001;10:231–241. [PubMed: 18249614]

Clarenz U, Diewald U, Rumpf M. Anisotropic geometric diffusion in surface processing. IEEE Visualization 2000 2000:397–405.

Demirkaya O. Anisotropic diffusion filtering of PET attenuation data to improve emission images. Phys. Med. Biol 2002;47:N271–N278. [PubMed: 12433127]

Elder JH, Zucker SW. Local scale control for edge detection and blur estimation. IEEE Trans. Pattern Anal. Machine Intell 1998;20:699–716.

Gerig G, Kübler O, Kikinis R, Jolesz FA. Nonlinear anisotropic filtering of MRI data. IEEE Trans. Med. Imaging 1992;11:221–232. [PubMed: 18218376]

Ghazel M, Freeman GH, Vrscay ER. Fractal image denoising. IEEE Trans. Image Process 2003;12:1560–1578. [PubMed: 18244711]

Gilboa G, Sochen N, Zeevi YY. A forward-and-backward diffusion process for adaptive image enhancement and denoising. IEEE Trans. Image Process 2002;11:689–703. [PubMed: 18244666]

Gilboa G, Sochen N, Zeevi Y. Image enhancement and denoising by complex diffusion processes. IEEE Trans. Pattern Anal. Machine Intell 2004;26:1020–1036.

He R, Narayana A. Automatic delineation of Gd enhancements on magnetic resonance images in multiple sclerosis. Med. Phys 2002;29:1536–1546. [PubMed: 12148736]

Höhne, K.; Fuchs, H.; Pizer, S. 3D Imaging in Medicine: Algorithms, Systems, Applications. Berlin, Germany: Springer-Verlag; 1990.

Kim HY, Giacomantone J, Cho ZH. Robust anisotropic diffusion to produce enhanced statistical parametric map from noisy fMRI. CVIU 2005;99:435–452.

Kim KK, Franz MO, Schölkopf B. Iterative kernel principal component analysis for image modeling. IEEE Trans. Pattern Anal. Machine Intell 2005;27:629–639.

Krissian K. Flux-based anisotropic diffusion applied to enhancement of 3-D angiogram. IEEE Trans. Med. Imaging 2002;21:1440–1442. [PubMed: 12575880]

Lee JS. Digital image enhancement and noise filtering by use of local statistics. IEEE Trans. Pattern Anal. Machine Intell 1980;PAMI-2:165–168.

Lester H, Arridge SR. A survey of hierarchical nonlinear medical image registration. Pattern Recognition 1999;32:129–149.

Lev A, Zucker SW, Rosenfeld A. Iterative enhancement of noisy images. IEEE Trans. Syst., Man, Cybern 1977;SMC-7:435–442.

Liang, P.; Wang, YF. Local scale controlled anisotropic diffusion with local noise estimate for image smoothing and edge detection; Proc. Int. Conf. Computer Vision; 1998. p. 193-200.

Madabhushi A, Udupa JK, Souza A. Generalized scale: theory, algorithms, properties and application to image inhomogeneity correction. Computer Vision Image Understanding 2006;101:100–121.

Meijering E, Niessen W, Weickert J, Viergever M. Diffusion-enhanced visualization and quantification of vascular anomalies in three-dimensional rotational angiography: Results of an in-vitro evaluation. Med. Image Anal 2002;6:215–233. [PubMed: 12270228]

Montagnat J, Sermesant M, Delingette H, Malandain G, Ayache N. Anisotropic filtering for model-based segmentation of 4D cylindrical echocardiographic images. Pattern Recognition Letters 2003;24:815–828.

Pal NR, Pal SK. A review of image segmentation techniques. Pattern Recognition 1993;26:277–1294.

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

Preuer T, Rumpf M. Anisotropic Nonlinear Diffusion in Flow Visualization. IEEE Visualization'99 1999:323–332.

Rosenfeld, A.; Kak, AC. Digital Picture Processing. Orlando, FL: Academic; 1982.

Rudin L, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms. Phys. D 1992;60:259–268.

Saha PK. Tensor scale: a local morphometric parameter with applications to computer vision and image processing. CVIU 2005;99:384–413.

Saha PK, Udupa JK. Optimum image thresholding via class uncertainty and region homogeneity. IEEE Trans. Pattern Anal. Machine Intell 2001;23:689–706.

Saha PK, Udupa JK. Scale-based diffusive image filtering preserving boundary sharpness and fine structures. IEEE Trans. Med. Imaging 2001;20:221–232.

Saha PK, Udupa JK, Odhner D. Scale-based fuzzy connected image segmentation: theory, algorithms, and validation. CVIU 2000;77:145–174.

Sapiro, G.; Tannenbaum, A. Edge preserving geometric smoothing of MRI data. Tech. Rep. Univ. of Minnesota, Dept. of Electrical Engineering; 1994.

Souza A, Udupa JK, Madabhushi A. Generalized scale-based image filtering. Proc. of SPIE: Medical Imaging 2005;5747:732–742.

Ter Haar Romeny, B., editor. Geometry-driven diffusion in computer vision. Dordrecht: Kluwer Academic Publishers; 1994.

Tschumperle D, Deriche R. Vector-Valued Image Regularization with PDE's: A Common Framework for Different Applications. IEEE Trans. Pattern Anal. Machine Intell 2005;27:506–517.

Udupa, JK.; Herman, G. 3D Imaging in Medicine. 2nd edition. Boca Raton, FL: CRC Press; 2000.

Udupa JK, Odhner D. Shell rendering. IEEE Comput. Graph. Applicat 1993;13:58–67.

Udupa JK, Odhner D, Samarasekera S, Goncalves R, Iyer K, Venugopal K, Furie S. 3DVIEWNIX: An open, transportable, multidimensional, multimodality, multiparametric imaging software system. Proc. of SPIE: Medical Imaging 1994;2164:58–73.

Udupa JK, Zhuge Y. Delineation operating characteristic (DOC) curve for assessing the accuracy behavior of segmentation algorithms. Proc. of SPIE: Medical Imaging 2004;5370:640–647.

Wang DCC, Vagucci AH, Li CC. Gradient inverse weighted scheme and the evaluation of its performance. Comput. Graph. Image Process 1981;15:167–181.

Wang X. On gradient inverse weighted filter (image processing). IEEE Trans. Signal Process 1992;40:482–484.

Weickert, J. Anisotropic Diffusion in Image Processing. Stuttgart: B. G. Teubner; 1998.

Weickert J. Coherence-Enhancing Diffusion Filtering. International Journal of Computer Vision 1999;31:111–127.

Wells WM, Viola P, Atsumi H, Nakajima S, Kikinis R. Multi-modal image registration by maximization of mutual information. Med. Image Anal 1996;1:35–51. [PubMed: 9873920]

Xu Y, Weaver JB, Healy DM Jr, Lu J. Wavelet transform domain filters: a spatially selective noise filtration technique. IEEE Trans. Image Process 1994;3:747–758. [PubMed: 18296244]

NIH-PA Author Manuscript

NIH-PA Author Manuscript

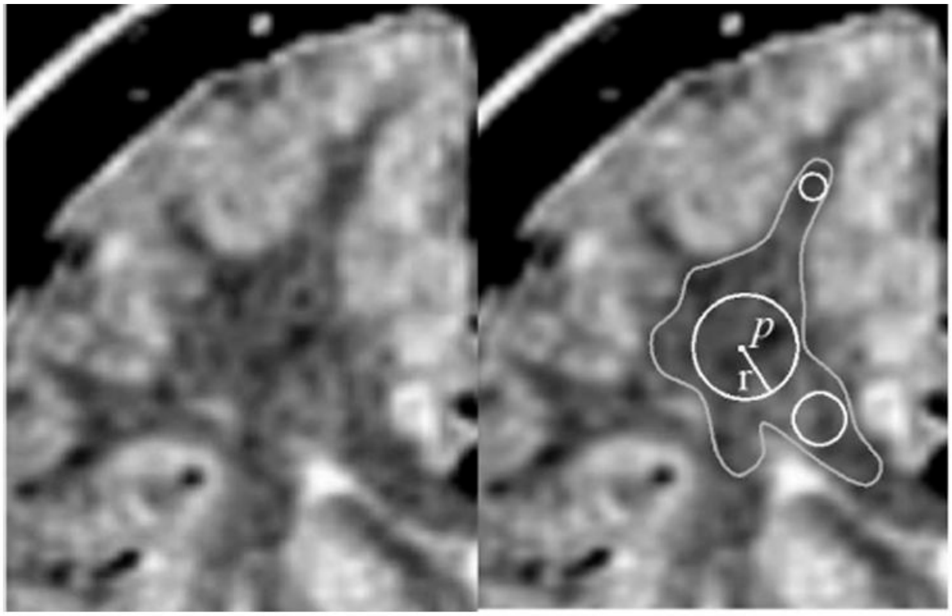

Figure 1.
Illustration of diffusive filtering control by local scale information. On left, a zoomed in axial PD MR image of the left superior portion of the head of a subject is shown. On right, the b-scale regions (white circles) and a g-scale region (light gray contour) of an image element *p* are overlaid. In b-scale-based diffusion, the radius *r* of the circle controls diffusion at *p*, while the size and shape of the g-scale region will control the diffusion process in g-scale-based diffusion.

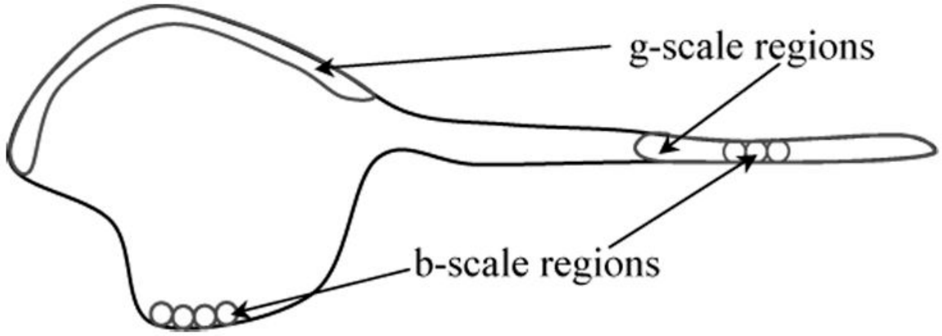

Figure 2.
An object region in an image consisting of a blobby and an elongated part is shown. The b-scale regions in the vicinity of the boundary of the object are individual small (radius=1) balls. The g-scale regions in the vicinity of the boundary of the object, however, are likely to be elongated thin regions running along the boundary. Thus, along the boundary, individual b-scales do not permit or encourage diffusion, whereas, since individual g-scale regions are likely to run along the boundary, they do encourage diffusion along boundaries. In the interior of the object, both b- and g-scales would behave in a similar manner.

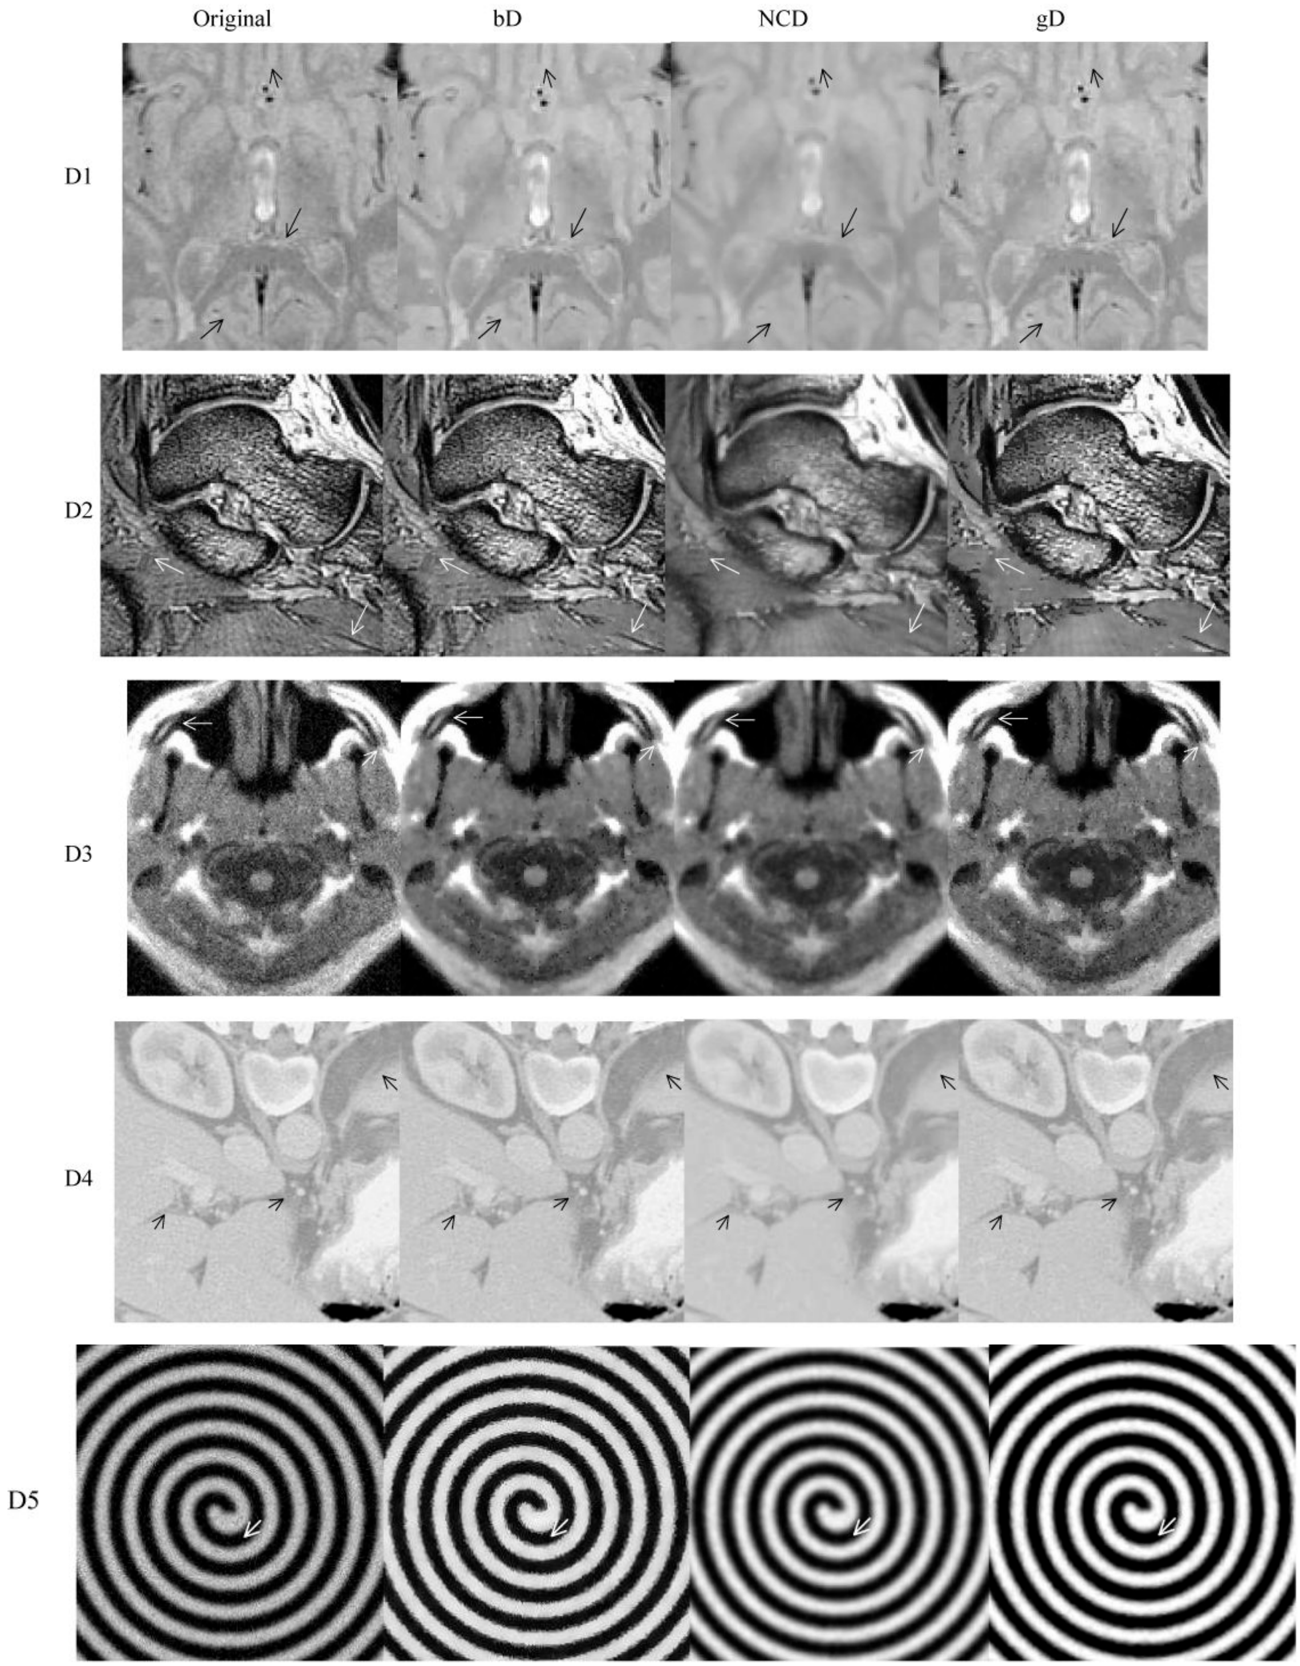
Original
bD
NCD
gD
D1
D2
D3
D4
D5

NIH-PA Author Manuscript
NIH-PA Author Manuscript

Figure 3.
Filtered scenes output by the three methods for data sets D1–D5.

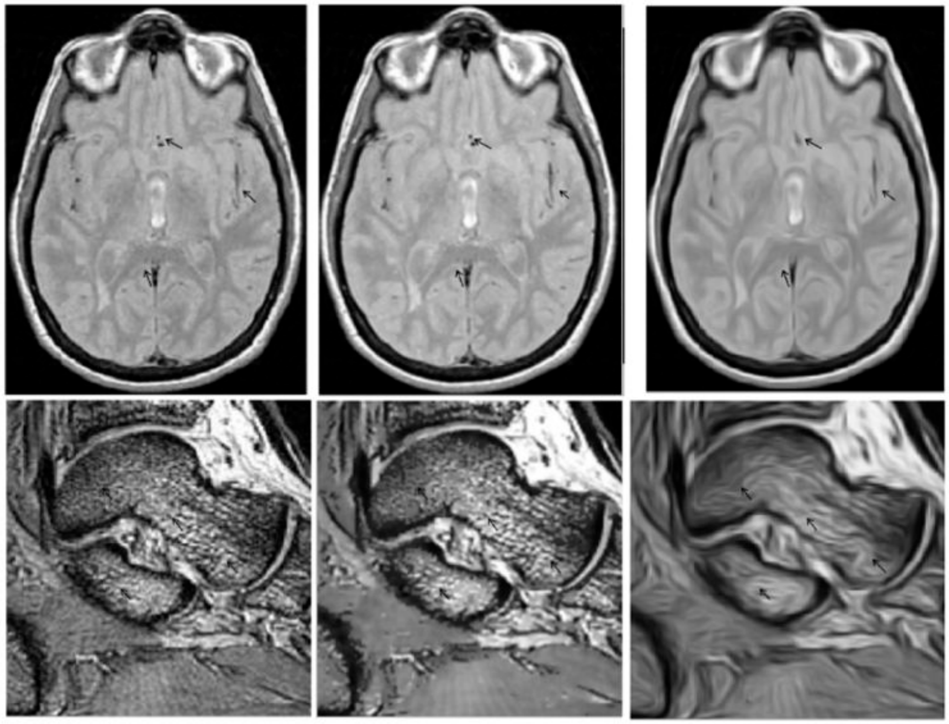

Figure 4.
Left to right: original, gD and Weickert et al. diffusion.

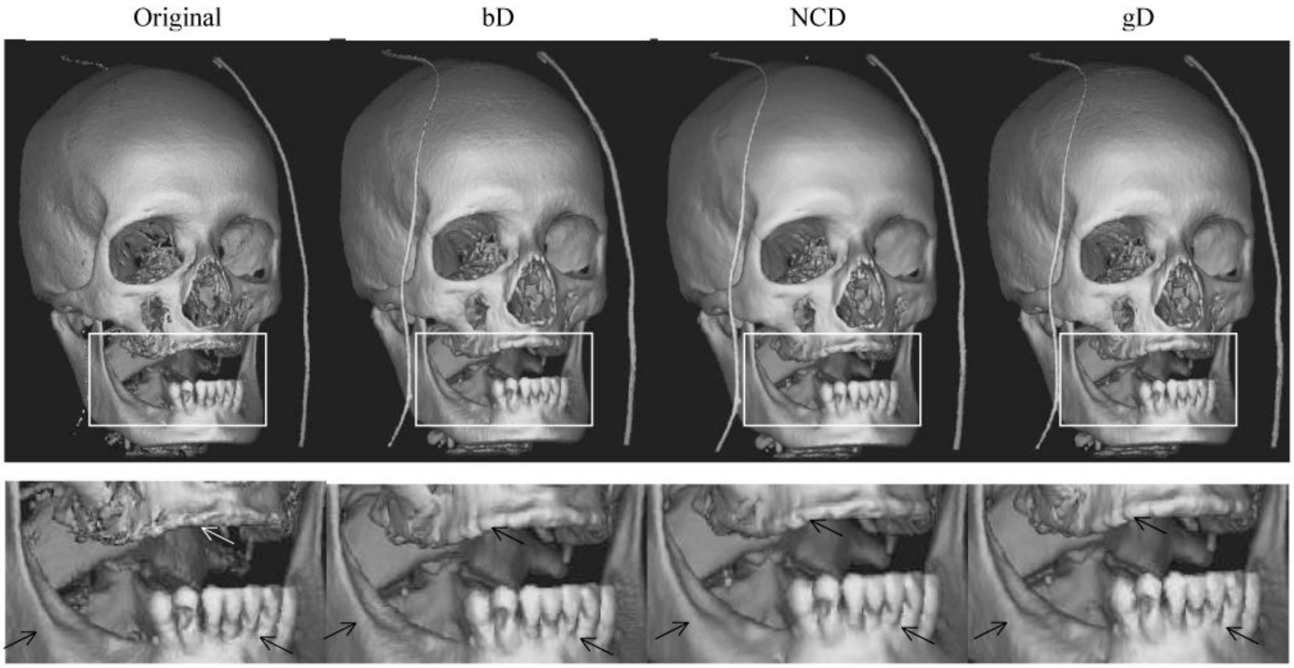

Figure 5.
Comparison of 3D renditions created from the filtered scenes for data set D6 to demonstrate the relative effectiveness of smoothing by the three methods along edges. Top row: Skull surfaces. Bottom row: Zoomed in views of the teeth and mandible.

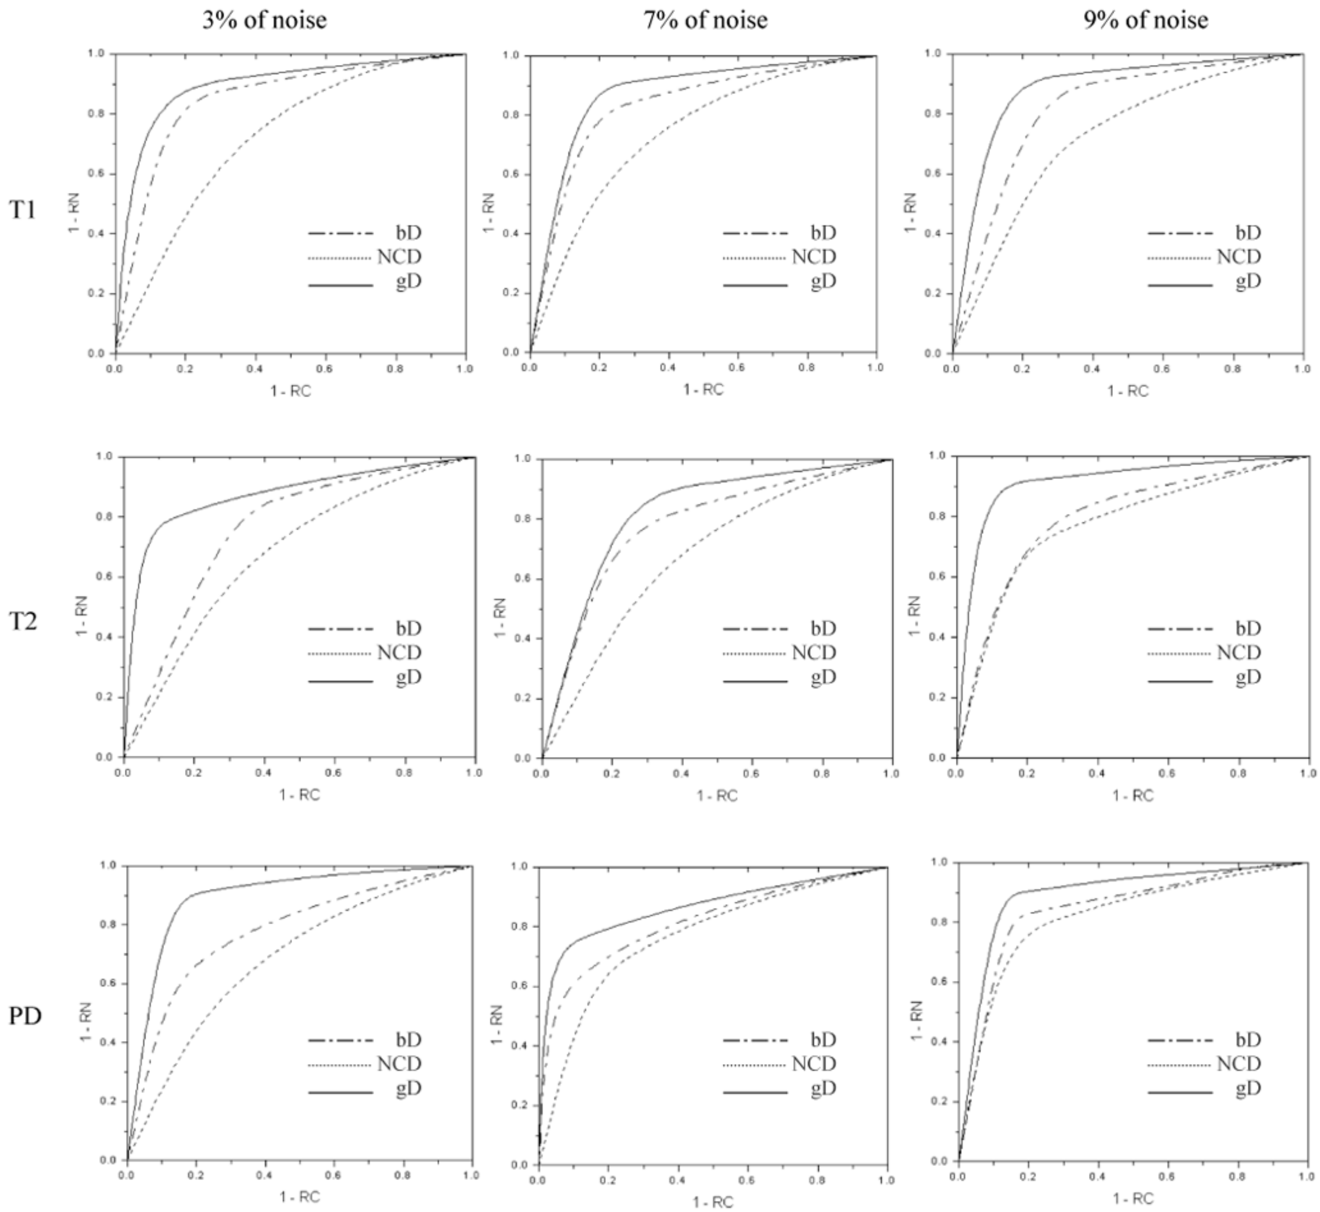

Figure 6.
Curves for 45 MR phantom scenes depicting the variation of *RN* and *RC* as the number of iterations is changed from 1 to 200.

Table 1

The data sets used in evaluation.

Data Set	Description	Voxel Size (mm^3)	Scene Domain
D1	PD brain MR image of a multiple sclerosis patient	0.86 × 0.86 × 3	256 × 256 × 51
D2	Gradient echo MR image of a normal human foot	0.55 × 0.55 × 1.3	256 × 256 × 60
D3	T1 simulated brain image with 9% noise from Brainweb (www.bic.mni.mcgill.ca/brainweb)	1 × 1 × 9	181 × 217 × 20
D4	CT image of a human thorax	0.78 × 0.78 × 2.5	512 × 512 × 60
D5	A 2D mathematical phantom constituting a thin Archimedes' spiral pattern	1×1	512×512
D6	CT image of a human head (Visible Woman)	0.49 × 0.49 × 1	512 × 512 × 209
D7	45 3D simulated MRI data sets of the brain from Brainweb containing PD, T1, T2 images at different levels of noise and slice thickness values	1×1× (1 – 9)	181 × 217 × (20 –181)

Table 2

AUC values for the curves shown in Figure 6.

Protocol	Noise level	bD	NCD	gD
T1	3%	0.8713	0.7289	0.9154
	7%	0.8596	0.7509	0.9061
	9%	0.8636	0.7551	0.9025
T2	3%	0.8100	0.6885	0.8837
	7%	0.8187	0.7811	0.8688
	9%	0.8335	0.7918	0.9311
PD	3%	0.7875	0.6908	0.9072
	7%	0.8166	0.7771	0.8601
	9%	0.8716	0.8398	0.8969

Table 3

AUC values observed for different values of σ_ψ on a T1 phantom scene.

	0.25 σ_ψ	**0.5 σ_ψ**	**0.75 σ_ψ**	**1.25 σ_ψ**	**1.5 σ_ψ**	**Mean AUC**	**SD AUC**
bD	0.8116	0.8302	0.8172	0.8042	0.7865	0.8099	0.0162
NCD	0.6767	0.7268	0.7270	0.7181	0.7148	0.7127	0.0208
gD	0.8772	0.8537	0.8032	0.8117	0.7928	0.8277	0.0361