

# MRI diffusion-based filtering: a note on performance characterisation

Ovidiu Ghita\*, Kevin Robinson, Michael Lynch, Paul F. Whelan

*Vision Systems Group, School of Electronic Engineering, Dublin City University, Glasnevin, Dublin 9, Ireland*

Received 30 August 2004; accepted 17 December 2004

## Abstract

Frequently MRI data is characterised by a relatively low signal to noise ratio (SNR) or contrast to noise ratio (CNR). When developing automated Computer Assisted Diagnostic (CAD) techniques the errors introduced by the image noise are not acceptable. Thus, to limit these errors, a solution is to filter the data in order to increase the SNR. More importantly, the image filtering technique should be able to reduce the level of noise, but not at the expense of feature preservation. In this paper we detail the implementation of a number of 3D diffusion-based filtering techniques and we analyse their performance when they are applied to a large collection of MR datasets of varying type and quality. © 2005 Elsevier Ltd. All rights reserved.

*Keywords:* MRI; Feature preserving smoothing; Anisotropic diffusion; Image segmentation

## 1. Introduction

In a survey on image smoothing techniques the approaches encountered may be classified under two broad headings, linear and non-linear [1,2]. Standard linear smoothing techniques based on local averaging or Gaussian weighted spatial operators reduce the level of noise, but this is achieved at the expense of poor feature preservation. Consequently, the filtered data appears blurry as step intensity discontinuities such as edges are attenuated. To compensate for these undesirable effects, non-linear techniques have been developed in order to achieve better feature preservation. Among these, the median filter is the simplest operator to remove impulse-like noise [2]. More complex non-linear techniques include statistical approaches based on non-parametric estimation [3,4]. However, while these methods do alleviate somewhat the shortcomings associated with linear techniques, they still perform only modestly when the data is affected by long tailed noise distributions. To complement these filtering approaches, a number of adaptive techniques have been proposed [2,3,5]. These methods try to achieve the best

trade-offs between smoothing efficiency, feature preservation and the generation of artefacts.

Recent developments based on non-linear diffusion [6–9] alleviate the major limitations associated with conventional linear and non-linear smoothing methods. Diffusion-based filtering was originally developed by Perona and Malik [10] in order to implement an optimal, feature preserving smoothing strategy. Many implementations follow their original approach where the main aim was to improve numerical stability [11,12]. This was advanced by Weikert [13] where he developed a new smoothing algorithm by permitting diffusion along the direction of edges. Gerig et al. [6] extended this work to 3D and evaluated its usefulness when applied to medical 2D and 3D datasets. In this paper our aim is to further extend their initial work in presenting the implementation of three diffusion-based smoothing algorithms where a special emphasis has been placed on performance characterisation [14,15]. We have evaluated the algorithms on various MRI datasets and the results are presented and discussed.

This paper is organised as follows. Section 2 describes the extension to 3D of the standard diffusion algorithm. Section 3 details the implementation of an adaptive 3D diffusion-based smoothing strategy. In Section 4 a 3D anisotropic Gaussian filtering technique is introduced. Section 5 presents an extensive performance characterisation of the smoothing strategies described in this paper and Section 6 presents some concluding remarks.

\* Corresponding author. Tel.: +353 1 700 7637; fax: +353 1 700 5508.  
E-mail address: ghitao@eeng.dcu.ie (O. Ghita).

## 2. Non-linear 3D diffusion filtering

In this section we describe the extension to 3D of the smoothing strategy that has been described in Perona and Malik's paper [10]. In their paper smoothing is formulated as a diffusive process where smoothing is performed at intra regions and suppressed at region boundaries.

This non-linear smoothing procedure can be defined in terms of the derivative of the flux function

$$u_t = \text{div}(D(|\nabla u|)\nabla u) \quad (1)$$

where  $u$  is the input data,  $D$  represents the diffusion function,  $t$  indicates the iteration step and  $\text{div}$  is the divergence operator. The smoothing strategy described in Eq. (1) can be translated into an iterative discrete formulation for 3D data as follows

$$I_{x,y,z}^{t+1} = I_{x,y,z}^t + \lambda \sum_{R=1}^6 [D(\nabla_R I) \nabla_R I]^t \quad (2)$$

where  $\nabla$  is the gradient operator than can be defined in a six voxel connected neighbourhood (see Eq. (3)) and  $\lambda$  is a contrast parameter that takes a value in the range  $0 < \lambda < 0.16$  as suggested in [10]. For a 26 voxel connected neighbourhood the implementation is similar and we just have to adjust the range for the  $\lambda$  parameter accordingly

$$\begin{aligned} \nabla_1 I_{x,y,z} &= I_{x-1,y,z} - I_{x,y,z}, & \nabla_2 I_{x,y,z} &= I_{x+1,y,z} - I_{x,y,z}, \\ \nabla_3 I_{x,y,z} &= I_{x,y-1,z} - I_{x,y,z}, & \nabla_4 I_{x,y,z} &= I_{x,y+1,z} - I_{x,y,z}, \\ \nabla_5 I_{x,y,z} &= I_{x,y,z-1} - I_{x,y,z}, & \nabla_6 I_{x,y,z} &= I_{x,y,z+1} - I_{x,y,z} \end{aligned} \quad (3)$$

The diffusion function  $D(x)$  should be bounded in the interval  $(0 \rightarrow 1)$  and should have the highest value when the input  $x$  has the value zero. These requirements translate to minimal smoothing around boundaries where the gradient has high values. In practice, a large number of functions can be engineered to satisfy these requirements and in our implementation we have used two types of diffusion functions, exponential and reciprocal

$$D(|\nabla I|) = e^{-(\nabla I/k)^2} \quad (4)$$

$$D(|\nabla I|) = \frac{1}{1 + \left(\frac{|\nabla I|}{k}\right)^2} \quad (5)$$

where  $k$  is the diffusion parameter. The parameter  $k$  controls the smoothing level, the smoothing being more pronounced for high values of  $k$ . The experimental data indicates that slightly better results are obtained when the exponential form illustrated in Eq. (4) is used in the expression depicted by Eq. (2).

## 3. Adaptive 3D diffusion smoothing

The adaptive smoothing algorithm implemented in this paper represents the extension to 3D of the algorithm proposed by Chen [16]. Chen [16] demonstrated that the standard diffusion algorithm might not return optimal results when applied to image data defined by a very low SNR. To tackle this limitation he proposed to use two discontinuity measures jointly in order to control the smoothing process. To this end, he used the spatial variance to measure contextual discontinuities and the gradient information as a measure of the local discontinuities.

In order to measure the local discontinuities we calculate the derivatives in a  $3 \times 3 \times 3$ , 13 diagonal derivatives calculated in a 26 voxel-connected neighbourhood

$$\begin{aligned} \nabla_1 I_{x,y,z} &= |I_{x+1,y,z} - I_{x-1,y,z}| \\ \nabla_2 I_{x,y,z} &= |I_{x,y+1,z} - I_{x,y-1,z}| \\ \nabla_3 I_{x,y,z} &= |I_{x+1,y+1,z} - I_{x-1,y-1,z}| \\ &\vdots \end{aligned} \quad (6)$$

where  $I_{x,y,z}$  represent the voxel intensity at the position  $(x,y,z)$ . Then, we can define the local discontinuity measure as:

$$E_{x,y,z} = \frac{\sum_{i=1}^{13} |\nabla_i|}{13} \quad (7)$$

As mentioned earlier the contextual discontinuities are sampled by measuring the spatial variance. The variance is computed in a cubic neighbourhood around the pixel of interest  $N_{xyz}(R)$  as follows

$$\sigma_{xyz}^2(R) = \frac{\sum_{(i,j,k) \in N_{xyz}(R)} (I_{i,j,k} - \mu_{xyz}(R))^2}{|N_{xyz}(R)|} \quad (8)$$

where  $\mu_{xyz}$  is the mean value of the voxels situated in the cubic neighbourhood  $N_{xyz}(R)$  and  $|N_{xyz}(R)|$  is the number of

Table 1

Performance characterisation results when the algorithms have been applied to an artificially created dataset

Noise type	SD noise	SD F1	SD F2	SD F3	Greyscale expected	Greyscale F1	Greyscale F2	Greyscale F3
G-15	13.72	1.91	1.62	2.06	127	127	128	128
G-30	31.93	7.64	3.03	5.57	127	128	129	133
P-15	13.02	1.07	0.76	1.74	127	139	138	138
P-30	26.97	9.60	7.62	3.69	127	141	141	142
W-15	4.63	1.50	0.21	0.69	127	126	127	127
W-30	8.56	1.71	0.60	1.14	127	125	126	127

$F1$ ,  $F2$ ,  $F3$  denote the standard diffusion, adaptive smoothing and anisotropic diffusion, respectively.

Table 2  
The RMS of the standard deviations of the homogenous areas for the original and filtered MRI datasets used in our experiments

	Heart	Brain	Whole body	MRCP
Original data	4.95	9.21	20.46	18.80
3D diffusion	1.88	6.28	14.47	10.96
3D adaptive	1.73	6.16	14.05	10.83
3D anisotropic	2.08	6.48	16.00	9.77

pixels. Then the spatial variances for the whole volume are scaled between 0 and 1 using the simple transformation illustrated in Eq. (9)

$$\tilde{\sigma}_{xyz}^2 = \frac{\sigma_{xyz}^2(R) - \sigma_{\min}^2(R)}{\sigma_{\max}^2(R) - \sigma_{\min}^2(R)} \quad (9)$$

The adaptive smoothing scheme is iterative and is illustrated below

$$I_{xyz}^{t+1} = I_{xyz}^t + \eta_{xyz} \frac{\sum_{i,j,k \in N_{xyz}(R) \setminus \{(x,y,z)\}} \eta_{ijk} \gamma_{ijk}^t (I_{ijk}^t - I_{xyz}^t)}{\sum_{i,j,k \in N_{xyz}(R) \setminus \{(x,y,z)\}} \eta_{ijk} \gamma_{ijk}^t} \quad (10)$$

where

$$\eta_{ijk} = \exp(-\alpha \Phi(\sigma_{ijk}^2(R), \theta_\sigma)) \quad (11)$$

$$\gamma_{ijk}^t = \exp(-E_{ijk}^t/S) \quad (12)$$

$$\Phi(\tilde{\sigma}_{xyz}^2(R), \theta_\sigma) = \begin{cases} 0 & \tilde{\sigma}_{xyz}^2(R) < \theta_\sigma \\ \tilde{\sigma}_{xyz}^2(R) & \tilde{\sigma}_{xyz}^2(R) \geq \theta_\sigma \end{cases} \quad (13)$$

In the formulation illustrated in Eq. (10) the parameters  $\theta_\sigma$  and  $\alpha$  control the extent to which the contextual

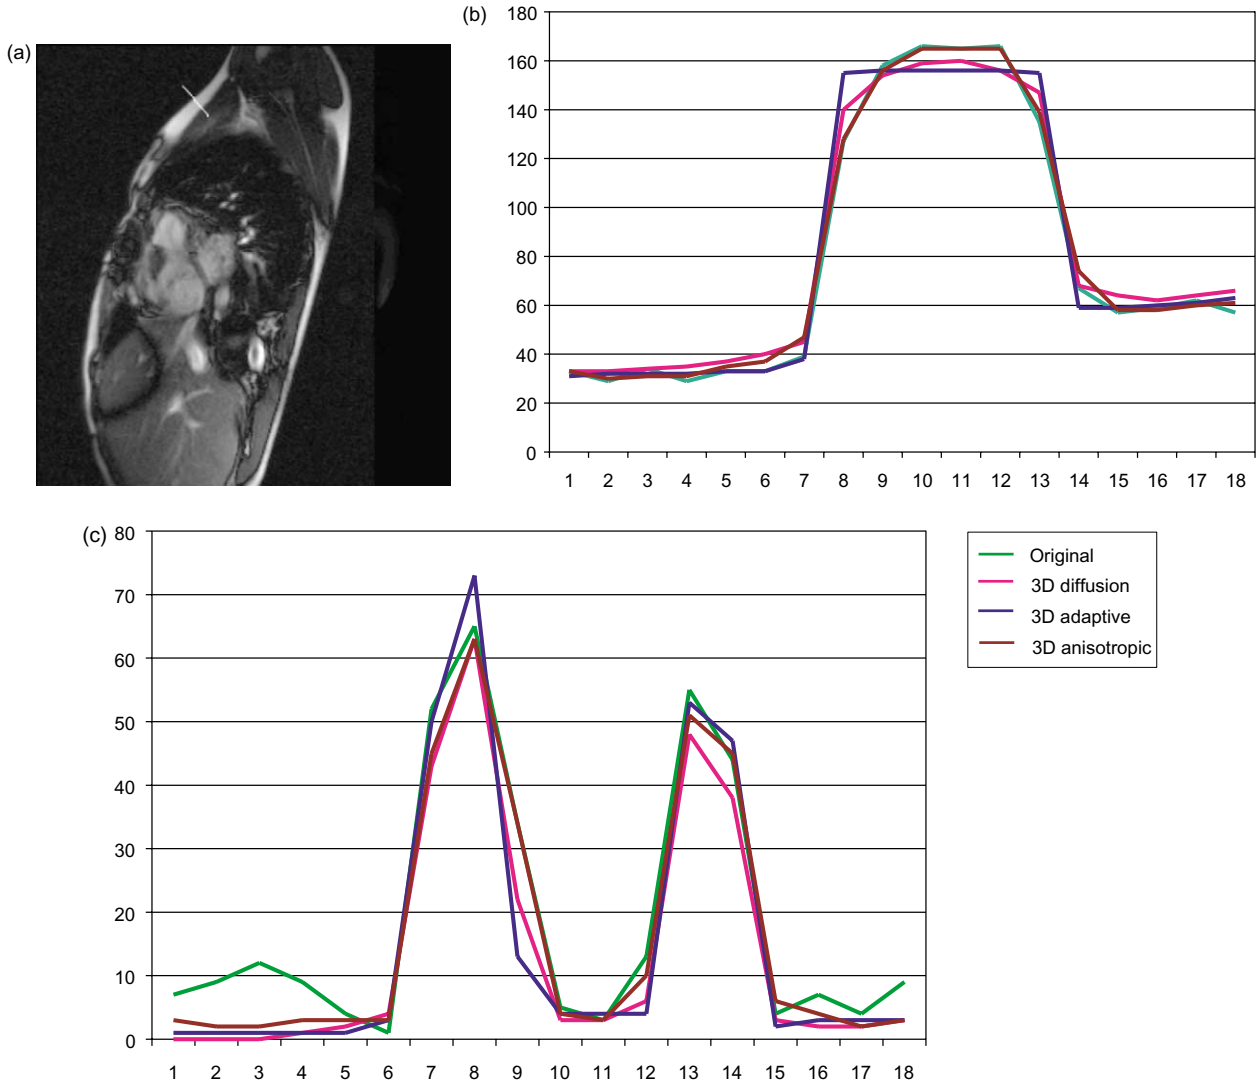


Fig. 1. (a) Slice of the heart MRI dataset. Pixel (b) and (c) gradient intensities are plotted for the highlighted edge illustrated in image (a).

discontinuities should be preserved while  $S$  is the diffusion parameter that controls the preservation of local discontinuities. It can be seen that  $\gamma$  implements the diffusion function depicted in Eq. (4). For this implementation the parameter  $R=3$  (for larger values image details such as thin lines are suppressed) and typical values for controlling parameters are  $\alpha=10$ ,  $\theta_\sigma=(0.05 \rightarrow 0.95)$ ,  $S=(5 \rightarrow 30)$ . The algorithm is typically run for 5–20 iterations.

#### 4. Anisotropic Gaussian smoothing

An anisotropic filter based on the familiar Gaussian model has also been implemented in order to provide edge enhancing, directional smoothing. This approach reduces to convolution with a scaled Gaussian mask where

the calculation of the kernel’s weights becomes the key issue governing the performance of this smoothing algorithm. By calculating the local gradient vector and favouring smoothing along the edge over smoothing across it, we can achieve boundary-preserving filtering where image regions are smoothed while edges are enhanced.

The weight for a neighbour pixel  $q$  can be calculated as a function of the gradient at point  $p$ , at the mask origin, and the distance from the origin to the neighbour  $q$ . The relationship used in our approach is given in Eq. (14) where  $\vec{pq}$  is the vector from the mask centre point  $p$  to the neighbour  $q$ ,  $\nabla u$  is the gradient vector at position  $p$ ,  $\lambda$  is the scale parameter that controls the smoothing strength while  $\mu$  is the shape parameter controlling anisotropy. It can be observed that in the case when  $\mu=0$  the anisotropic term  $((\vec{pq} \cdot \nabla u)/\lambda)^2(2\mu + \mu^2)$  is also zero and the smoothing

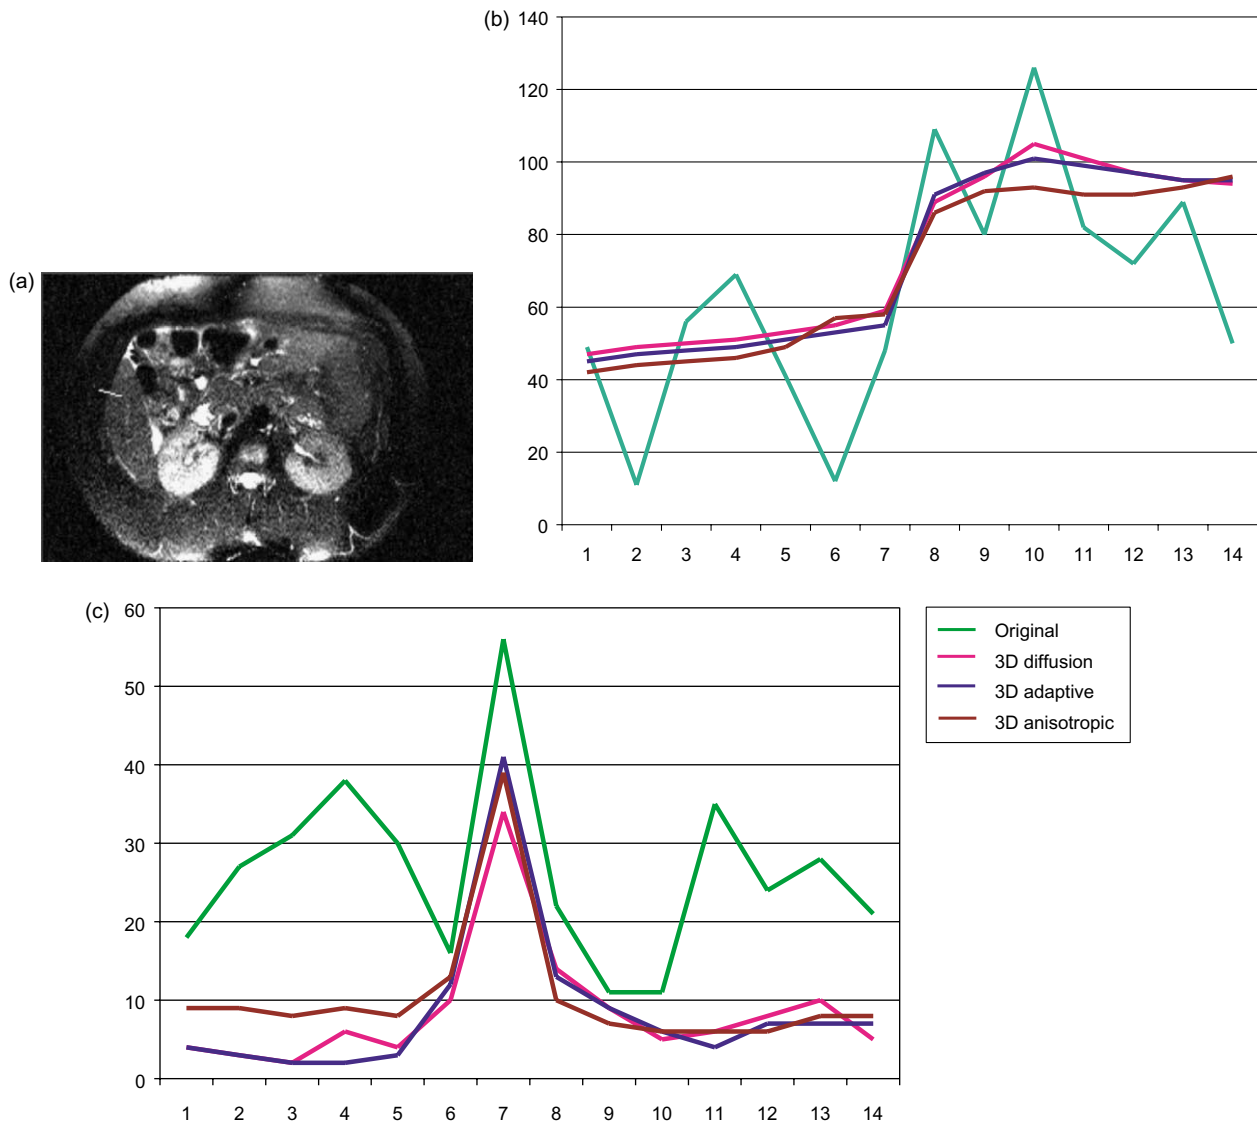


Fig. 2. (a) Slice of the MRCP dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge illustrated in image (a).

operation reduces to the non-linear isotropic form where smoothing is suppressed at image boundaries (as no directionality is applied)

$$wt(pq, \nabla u) = \exp\left(-\left(\frac{|\overline{pq}||\nabla u|}{\lambda}\right)^2 + \left(\frac{pq \cdot \nabla u}{\lambda}\right)^2 (2\mu + \mu^2)\right) \tag{14}$$

The smoothing operation filters the data iteratively with the 3D kernel constructed from Eq. (14). As the number of iterations is increased, more noise and small features are eliminated, but even in extreme cases the strong edges in the data are well preserved in both location and strength.

**5. Experiments and results**

The filtering algorithms described in this paper have been applied to a large number of MR datasets including MRCP, whole body, brain and heart sequences. The aim of these

experiments is to conduct a detailed performance characterisation for smoothing algorithms described in this paper in order to produce quantitative results. In our experiments we also included datasets where additional noise had been added to the original sequences.

To evaluate the performance of the smoothing algorithms described in this paper, the first set of experiments were conducted on a synthetic dataset that is defined by a homogenous cubic object with a known greyscale value surrounded by background pixels. To test smoothing algorithms on this artificial dataset is advantageous as the ground truth data is known and the smoothing results are easy to evaluate. We tested the efficiency of the algorithms when the artificial dataset was corrupted with various types of 3D image noise, including Gaussian, Poisson and additive uniformly distributed white noise [1,17]. As quantitative values we have evaluated the local uniformity sampled by the 7×7×7 standard deviation at the location situated at the centre of the cube and the alteration of the greyscale value at the same position when compared with

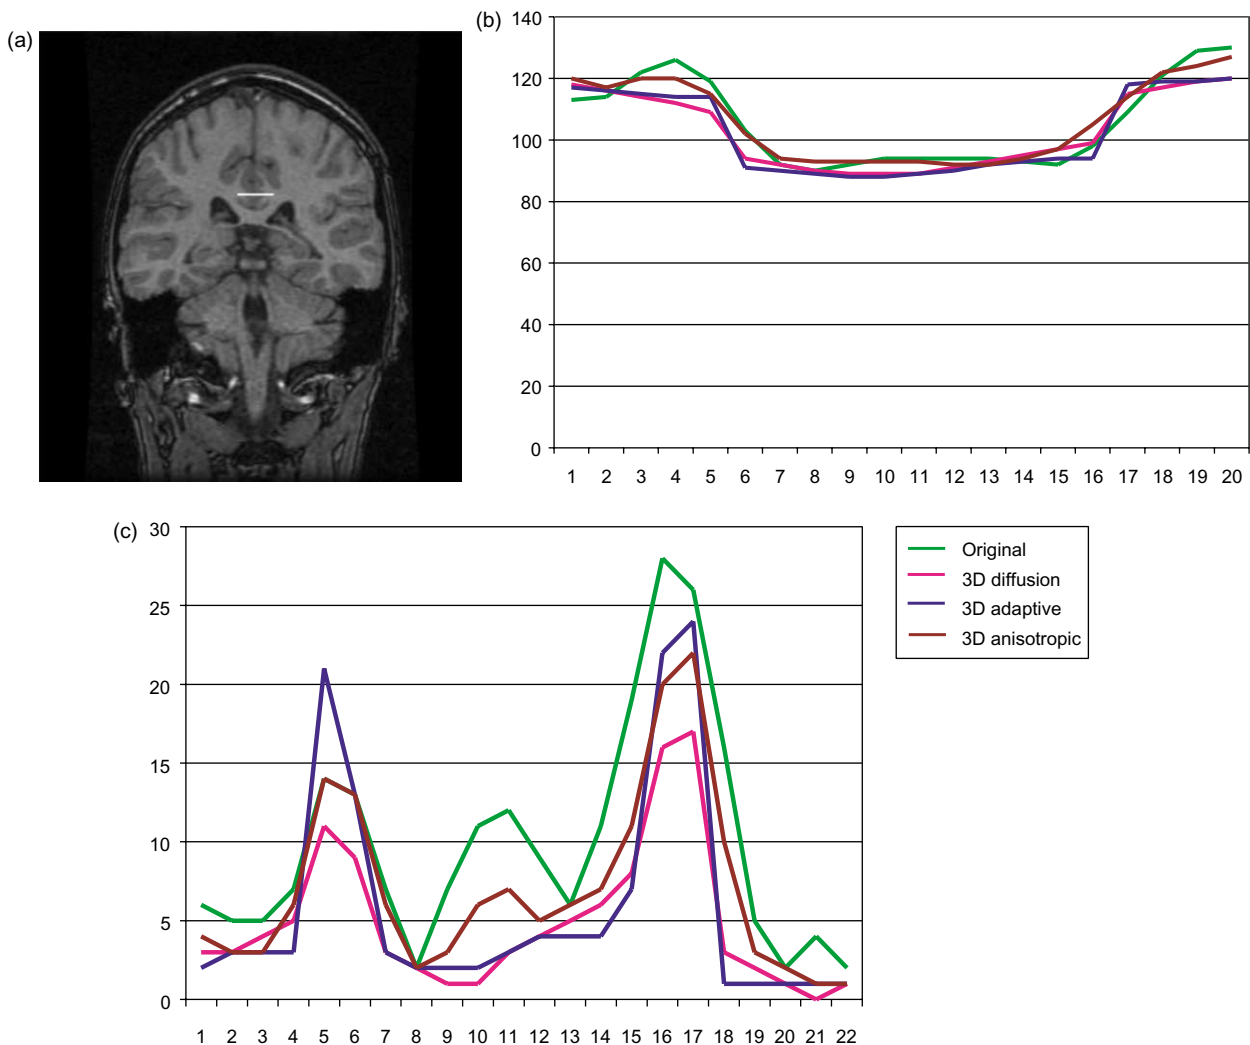


Fig. 3. (a) Slice of the brain MRI dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge in image (a).

the expected known value. Some experimental results are depicted in Table 1.

In Table 1 the symbols G-15 and G-30 indicate that the synthetic dataset has been corrupted with Gaussian noise (standard deviation 15 and 30 greyscale values). Similarly P-15 and P-30 denote the fact that the test dataset has been corrupted with Poisson noise (distribution 15 and 30 greyscale values) and W-15 and W-30 indicate that the dataset has been corrupted with uniformly distributed white noise (mean deviation 15 and 30 greyscale values).

In order to evaluate globally the noise removal efficiency on real datasets we need to define quantitative measures that indicate the overall performance of the smoothing algorithms that are evaluated. In this regard, we propose to evaluate jointly two quantitative measurements: the smoothness factor that assesses the global uniformity and the edge preservation factor that indicates to what extent the strong edge features are retained and enhanced. To this end, we employed the standard deviation as

a measure to evaluate the image local homogeneity. To be statistically relevant [17] the standard deviation should be calculated over a large region, but on the other hand the results will be affected by small non-uniformities such as intensity gradients or structural image variations [6]. This requirement is quite difficult to be accomplished if we want to develop an automatic performance characterisation scheme where user intervention is not required. One solution has been advanced by Canny [18] when he decided to select the threshold parameters for an edge detector based on analysis of the cumulative histogram of the gradients. However, due to the nature of MR datasets this criteria to identify the gradients generated by noise proved to be inefficient. Thus, in our implementation we have developed an alternative strategy based on observation. In this sense, we computed the standard deviation for all voxels in the original dataset in a  $7 \times 7 \times 7$  neighbourhood. These values were sorted with respect to their magnitude and from these values the 25% of

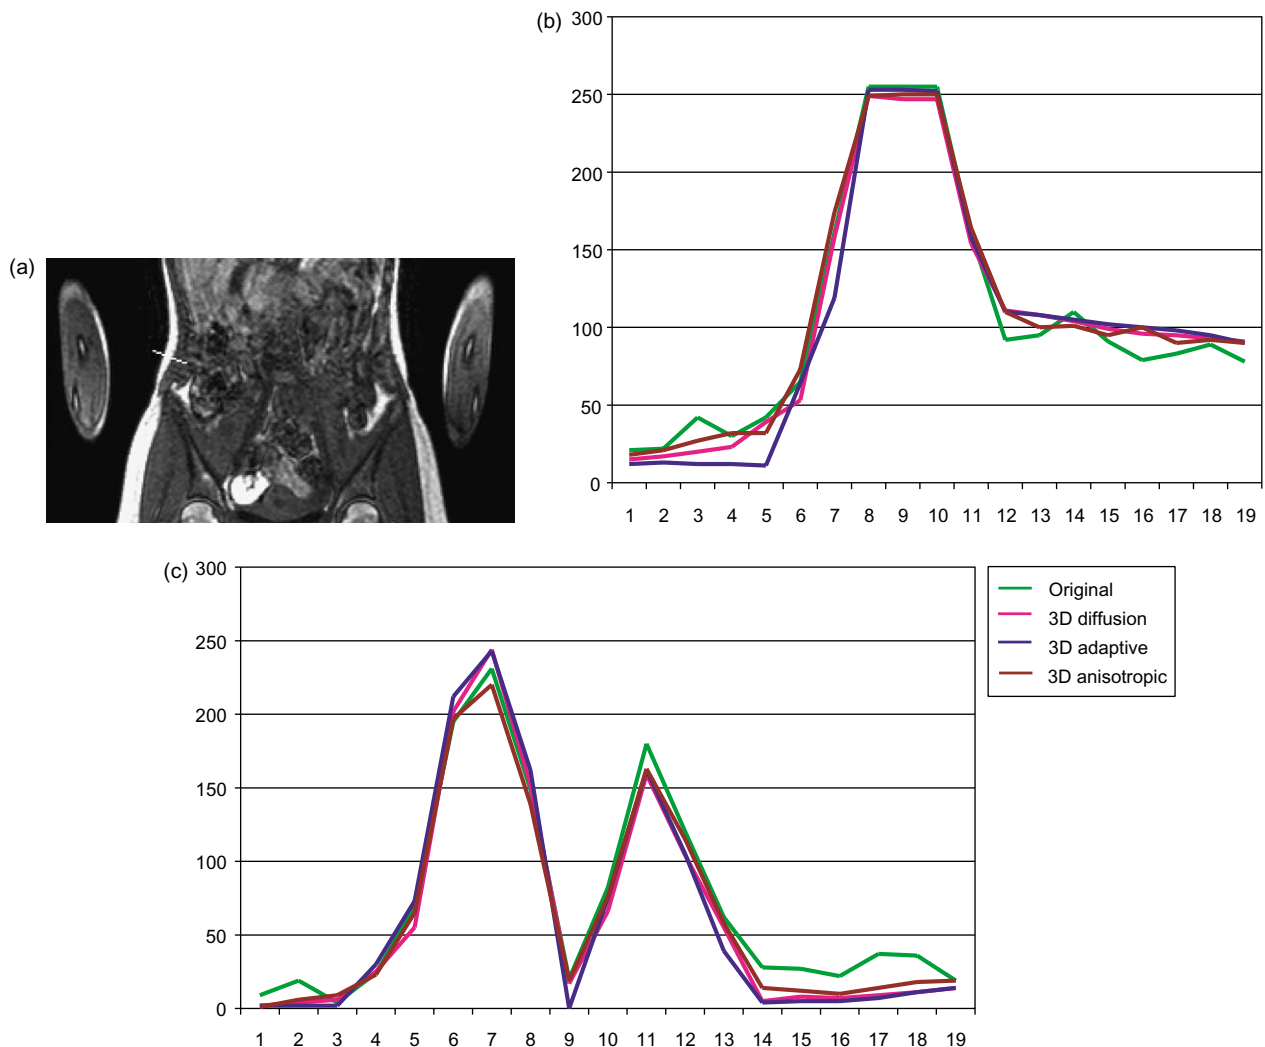


Fig. 4. (a) Slice of the whole body MRI dataset. Pixel (b) and gradient intensities (c) are plotted for the highlighted edge illustrated in image (a).

the highest values were eliminated, as they are likely to belong to edges and 25% of the lowest values are also eliminated as they are calculated from areas that have no significant texture (such as image regions defined by air). This strategy was applied to select the seed points that

belong to image regions defined by a low SNR. Then, the standard deviation for each of the filtered datasets is measured at the same voxel locations (also in a  $7 \times 7 \times 7$  neighbourhood). To evaluate a quantitative estimation we calculate the RMS value of the standard deviations from




Fig. 5. 3D data clustering results—heart dataset. (First row) Original dataset (slice 9) and corresponding image resulted after clustering. (Second row) 3D diffusion smoothed data (slice 9) and corresponding image resulted after clustering. (Third row) 3D adaptive smoothed data (slice 9) and corresponding image resulted after clustering. (Forth row) 3D anisotropic smoothed data (slice 9) and corresponding image resulting after clustering.




Fig. 6. 3D data clustering results—brain dataset. (First row) Original dataset (slice 4) and corresponding image resulted after clustering. (Second row) 3D diffusion smoothed data (slice 4) and corresponding image resulted after clustering. (Third row) 3D adaptive smoothed data (slice 4) and corresponding image resulted after clustering. (Forth row) 3D anisotropic smoothed data (slice 4) and corresponding image resulting after clustering.




Fig. 7. 3D data clustering results—MRCP dataset. (First row) Original dataset (slice 10) and corresponding image resulted after clustering. (Second row) 3D diffusion smoothed data (slice 10) and corresponding image resulted after clustering. (Third row) 3D adaptive smoothed data (slice 10) and corresponding image resulted after clustering. (Fourth row) 3D anisotropic smoothed data (slice 10) and corresponding image resulting after clustering.

the original and smoothed datasets resulting after the application of the smoothing strategies described in previous sections (for details refer to Table 2).

The edge strength is evaluated by plotting the intensity and gradient data at selected locations where

edges are located, before and after the application of the smoothing operations. Some graphical results are depicted in Figs. 1–4. The experimental data illustrated in Figs. 1–4 indicate that the 3D adaptive smoothing and 3D anisotropic smoothing algorithms perform better




Fig. 8. 3D data clustering results—whole body dataset. (First row) Original dataset (slice 6) and corresponding image resulted after clustering. (Second row) 3D diffusion smoothed data (slice 6) and corresponding image resulted after clustering. (Third row) 3D adaptive smoothed data (slice 6) and corresponding image resulted after clustering. (Forth row) 3D anisotropic smoothed data (slice 6) and corresponding image resulted after clustering.

than the standard diffusion. The 3D adaptive smoothing algorithm returned better results than the 3D anisotropic when applied to heart, brain and whole body datasets. The 3D anisotropic algorithm performed better when applied to MRCP dataset.

The graphs illustrated in Figs. 1 and 4 demonstrate the edge enhancement around image data defined by step-like edges. It can be noticed that the edge localisation is significantly improved. The effect of edge strengthening is even more pronounced for weaker edges in an MRI brain sequence (see Fig. 3) or in image areas affected by a high level of noise, as is the case of the MRCP dataset illustrated in Fig. 2.

The performance of the smoothing algorithms described in this paper is remarkable in discriminating a true edge from image noise (see Fig. 2c). Also notice the improved performance of the adaptive 3D smoothing algorithm as compared with the performance of the standard diffusion and the 3D anisotropic diffusion algorithms.

In order to emphasise the effectiveness of the smoothing strategies described in this paper we also present the segmentation resulting after the application of a 3D clustering algorithm [19] to the original and smoothed data. Samples of the segmentation results are depicted in Figs. 5–8.

## 6. Conclusions

In this paper we have described the implementation of three diffusion-based smoothing schemes and their application to medical 3D data. Our interest has focused on MRI acquisition modalities as MRI datasets are characteristically defined by a low signal to noise ratio (SNR). Hence, our aim was to demonstrate that far superior results are achieved if the MRI data is initially filtered in order to reduce the level of image noise and improve the SNR. In this regard, we have performed a detailed performance characterisation for each smoothing operator evaluated in this paper on both synthetic and real data (including heart, brain, whole body and MRCP image sequences). We conclude that the diffusion-based smoothing techniques offer an efficient approach to noise reduction, and more important this advantage is not achieved at the expense of feature preservation. The experimental data presented and discussed in this paper highlights the ability of the diffusion-based smoothing schemes to distinguish the high gradient image features from the MRI image acquisition noise.

The source code for the 3D smoothing schemes presented in this paper can be downloaded from the following web page: <http://www.eeng.dcu.ie/~whelanp/vsg/vsgcode.html>.

## References

- [1] Orfanidis SJ. Introduction to signal processing. Englewood Cliffs, NJ: Prentice-Hall; 1996.
- [2] Sonka M, Hlavac V, Boyle R. Image processing, analysis and machine vision. 2nd ed. Boston: PWS; 1998.
- [3] Smolka B, Lukac R, Chydzinski A, Plataniotis KN, Wojciechowski K. Fast adaptive similarity based impulse noise reduction filter. *Real Time Imaging* 2003;9(4):261–76.
- [4] Tang K, Astola J, Neuovo Y. Nonlinear multivariate image filtering techniques. *IEEE Trans. Image Process.* 1995;4(6):788–97.
- [5] Comaniciu D, Meer P. Mean shift: a robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Machine Intell.* 2002;24(5):603–19.
- [6] Gerig G, Kikinis R, Kubler O, Jolesz FA. Nonlinear anisotropic filtering of MRI data. *IEEE Trans. Med. Imaging* 1992; 11(2):221–32.
- [7] Montagnat J, Sermesant M, Delingette H, Malandain G, Ayache N. Anisotropic filtering for model-based segmentation of 4D cylindrical echocardiographic images. *Pattern Recognit. Lett.* 2003; 24(4/5):815–28 [Special Issue on Ultrasonic Image Processing and Analysis].
- [8] Sanchez-Ortiz GI, Rueckert D, Burger P. Knowledge-based tensor anisotropic diffusion of cardiac magnetic resonance images. *Med. Image Anal.* 1999;3(1):77–101.
- [9] Whitaker RT, Pizer SM. A multi-scale approach to nonuniform diffusion. *CVGIP: Image Understanding* 1993;57(1):111–20.
- [10] Perona P, Malik J. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Machine Intell.* 1990;12(7):629–39.
- [11] Nordstrom N. Biased anisotropic diffusion. A unified regularization and diffusion approach to edge detection. *Image Vis. Comput.* 1990; 8(4):318–27.
- [12] Weickert J, ter Haar Romeny BM, Viergever MA. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Trans. Image Process.* 1998;7(3):398–410.
- [13] Weickert J. Anisotropic diffusion in image processing. Stuttgart: Teubner Verlag; 1998.
- [14] Foerstner W. 10 Pros and cons against performance characterisation of vision algorithms. Proceedings of ECCV workshop on performance characteristics of vision algorithms, Cambridge, UK 1996.
- [15] Haralick RM. Performance characterization in computer vision. *CVGIP: Image Understanding* 1994;60(2):245–9.
- [16] Chen K. A feature preserving adaptive smoothing method for early vision. Technical report. Beijing, China: National Laboratory of Machine Perception, Centre for Information Science, Peking University; 1999.
- [17] Duda RO, Hart PE, Stork DG. Pattern classification. 2nd ed. New York: Wiley; 2001.
- [18] Canny J. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Machine Intell.* 1986;8(6):679–98.
- [19] Lynch M, Ghita O, Whelan PF. Comparison of 2D and 3D clustering on short axis magnetic resonance images of the left ventricle. CARS 2004 computer assisted radiology and surgery, June 23–26, Chicago, USA 2004.

**Ovidiu Ghita** received his BE and ME degrees in Electrical Engineering from Transilvania University, Brasov, Romania. From 1994 to 1996 he was an Assistant Lecturer in the Department of Electrical Engineering at Transilvania University. Since then, he has been a member of the Vision Systems Group at Dublin City University (DCU) during which time he received his PhD for work in the area of robotic vision. Currently, he holds a position of Postdoctoral Research Assistant in the Vision Systems Laboratory at DCU. His current research interests are in the area of range acquisition, shape representation, machine vision and medical imaging.

**Kevin Robinson** received his BE in Electronic Engineering from the University of Exeter in 1993. He graduated Master of Science in Computation from the University of Oxford, Programming Research Group a year later. He then embarked on a career in industry holding various positions in research and development over a period of 6 years, working in medical imaging, computer vision, and industrial inspection. During this time he attained chartered engineer status (CE) and corporate membership of the Institution of Electrical Engineers (MIEE). At the end of 2001 he commenced work on his PhD with the Vision Systems Group at Dublin City University applying advanced image processing techniques to complex problems in the field of medical image analysis and visualisation.

**Michael Lynch** received his BE (Hons) degree in Mechatronic Engineering from Dublin City University, Dublin, Ireland in 2002. He joined the Vision Systems Group in Dublin City University at the start of 2003 as a research student working towards an PhD. Current research interests include image processing and medical imaging, with emphasis on cardiac imaging.

**Paul F Whelan** received his BE (Hons) degree from the National Institute for Higher Education Dublin, a ME degree from the University of Limerick, and his PhD from the University of Wales, Cardiff. During the period 1985–1990 he was employed by Industrial and Scientific Imaging Ltd and later Westinghouse (WESL), where he was involved in the research and development of industrial vision systems. He was appointed to the School of Electronic Engineering, Dublin City University (DCU) in 1990 and currently holds the position of Associate Professor and Director of the Vision Systems Laboratory. As well as a wide range of scientific publications, Prof. Whelan co-edited *Selected Papers on Industrial Machine Vision Systems* (1994), and was the co-author of *Intelligent Vision Systems for Industry* (1997) and *Machine Vision Algorithms in Java* (2000). His research interests include applied morphology, texture analysis, machine vision and medical imaging. He is a Senior Member of the IEEE, a Chartered Engineer and a member of the IEE, SPIE and IAPR. He is also a member of a number of machine vision related conference program committees. He currently serves on the IEE Irish centre committee, as member of the governing board of the International Association for Pattern Recognition (IAPR) and as the President of the Irish Pattern Recognition and Classification Society.