

# Knowledge-based tensor anisotropic diffusion of cardiac magnetic resonance images

Gerardo I. Sanchez-Ortiz<sup>1,2\*</sup>, Daniel Rueckert<sup>1,3</sup> and Peter Burger<sup>1</sup>

<sup>1</sup>Department of Computing, Imperial College, London, UK

<sup>2</sup>Present address: Medical Vision Laboratory, Robotics Research Group, Department of Engineering Science, University of Oxford, Oxford, OX1 3JP, UK

<sup>3</sup>Present address: Computational Imaging Science Group, Division of Radiological Sciences and Medical Engineering, Guy's, King's and St Thomas' School of Medicine, King's College London, Guy's Hospital, London SE1 9RT, UK

## Abstract

We present a general formulation for a new knowledge-based approach to anisotropic diffusion of multi-valued and multi-dimensional images, with an illustrative application for the enhancement and segmentation of cardiac magnetic resonance (MR) images. In the proposed method all available information is incorporated through a new definition of the conductance function which differs from previous approaches in two aspects. First, we model the conductance as an explicit function of time and position, and not only of the differential structure of the image data. Inherent properties of the system (such as geometrical features or non-homogeneous data sampling) can therefore be taken into account by allowing the conductance function to vary depending on the location in the spatial and temporal coordinate space. Secondly, by defining the conductance as a second-rank tensor, the non-homogeneous diffusion equation gains a truly anisotropic character which is essential to emulate and handle certain aspects of complex data systems. The method presented is suitable for image enhancement and segmentation of single- or multi-valued images. We demonstrate the efficiency of the proposed framework by applying it to anatomical and velocity-encoded cine volumetric (4-D) MR images of the left ventricle. Spatial and temporal *a priori* knowledge about the shape and dynamics of the heart is incorporated into the diffusion process. We compare our results to those obtained with other diffusion schemes and exhibit the improvement in regions of the image with low contrast and low signal-to-noise ratio.

**Keywords:** anisotropic diffusion, geometry-driven diffusion, left-ventricle modelling, phase-sensitive MR, segmentation

*Received October 29, 1997; revised June 12, 1998; accepted November 19, 1998*

## 1. INTRODUCTION

Image processing and computer vision have traditionally dealt with problems such as image segmentation (i.e. dividing an image into meaningful and disjoint regions) in cases where the data can be expressed as a single- or vector-valued image function defined on an  $n$ -dimensional image domain. Many semantic interpretations of these image

functions rely on the extraction of geometric features such as edges, corners or ridges (Lindeberg and ter Haar Romeny, 1994). Determining at which scale of resolution these image features should be measured has emerged as a fundamental problem especially in cases where the image is affected by noise or any type of spurious artefacts that introduce unwanted variations of the image intensity. In this section we briefly review methods for generating images at different scales of resolution and explain how these schemes can be modified to achieve a more meaningful feature enhancement and image segmentation.

\*Corresponding author  
(e-mail: giso@robots.ox.ac.uk)

### 1.1. Linear diffusion

In recent years many different approaches for image descriptions based on linear and non-linear diffusion processes have been developed [a detailed review can be found in ter Haar Romeny (1994) and Weickert (1997)]. This concept originally emerged from the idea of analysing images at varying levels of resolution (in a scale space) and was first developed by Witkin (1983) and Koenderink (1984).

Koenderink pointed out that blurring can be expressed in terms of the heat conduction or diffusion equation:

$$\frac{\partial I(x, y, \tau)}{\partial \tau} = \nabla \cdot c \nabla I(x, y, \tau), \quad (1)$$

where  $\nabla$  is the gradient and  $\nabla \cdot$  is the divergence operator. If  $c$  is a constant, the diffusion process is called linear diffusion and the Gaussian kernel

$$G(x, y, \sigma) = \frac{1}{(2\pi\sigma^2)} e^{-(x^2+y^2)/2\sigma^2} \quad (2)$$

is the Green's function of Equation (1). In this case the diffused image can be obtained directly by a spatial convolution of the original image with the Gaussian function, i.e.

$$I(x, y, \tau) = I(x, y) \otimes G(x, y, \sigma), \quad (3)$$

and the scale parameter replaces the time parameter in Equation (1) with  $\sigma = \sqrt{2\tau c}$  given the initial condition  $I(x, y, 0) = I(x, y)$ . Moreover, the conductance term is often chosen to be  $c = 1$  so that it does not appear explicitly in the diffusion equation.

### 1.2. Non-linear diffusion

A scale space generated by a linear diffusion process as defined in Equation (1) causes an isotropic flow of intensities and is space-invariant. The resulting blurring of the image, however, eliminates not only noise but also the underlying signal which one ultimately wants to characterize. The result is a blurring process which eliminates successively more and more details of the image without taking the local image features into account. This leads to two problems: first, the blurring is coupled with a loss of information about image features such as small objects and sharp boundaries; secondly, the blurring tends to degrade the localization of the remaining image features.

#### 1.2.1. Edge-based diffusion

Based on this observation, Perona and Malik (1990) proposed a non-linear diffusion scheme in which the constant conductance term is replaced by a variable conductance term which can vary over space and time:

$$\frac{\partial I(x, y, \tau)}{\partial \tau} = \nabla \cdot c(x, y, \tau) \nabla I(x, y, \tau). \quad (4)$$

Perona and Malik suggested a conductance term which is a monotonically decreasing function of the magnitude of the gradient of the intensity:

$$c(x, y, \tau) = g(\|\nabla I\|) = \frac{1}{1 + (\|\nabla I\|/k)^2}. \quad (5)$$

This function introduces a new parameter  $k$  which controls the influence of the gradient. In practice  $k$  acts as a threshold which determines whether to preserve edges or not: areas in which the gradient magnitude is lower than  $k$  will be blurred more strongly than areas with a higher gradient magnitude. This tends to smooth uniform regions, while preserving the edges between different regions.

Such a diffusion process, Catté *et al.* (1992) have argued, is of a recursive nature and in cases where the image is degraded by noise can introduce uncorrelated as well as unbounded gradients. This makes it impossible to distinguish between edges which should be preserved and noise which should be diffused. Moreover, Catté *et al.* have argued that the diffusion process is not well posed because  $f \cdot g(f)$  is not monotonic. In practice, this can lead to significantly different results for very similar images. These problems can be avoided if the calculation of the gradient is regularized by a convolution with a Gaussian kernel:

$$\frac{\partial I(x, y, \tau)}{\partial \tau} = \nabla \cdot g(\|\nabla G(x, y, \sigma) \otimes I(x, y, \tau)\|) \nabla I. \quad (6)$$

Gerig and Whitaker (Gerig *et al.*, 1992; Whitaker and Gerig, 1994) have extended this approach to three-dimensional (3-D) images as well as to multi-valued  $T_1$  and  $T_2$  weighted magnetic resonance (MR) images.

Equations (4) and (5) and their regularized version correspond to (non-linear) inhomogeneous isotropic diffusion. However, the discretization used by Perona and Malik gives the scheme an anisotropic character, which will be discussed in Subsubsection 1.2.3. Since this scheme has been widely referred to as anisotropic diffusion we will refer to its implementation as the standard anisotropic diffusion method.

#### 1.2.2. Geometry-based diffusion

The non-linear diffusion schemes discussed so far are characterized by the evolution of the intensity function. Alternatively, one can characterize the diffusion as an evolution of the isophotes curves of the image. In fact, Niessen *et al.* (1997) have shown that both approaches are dual in the sense that one determines the other. One geometry-based diffusion scheme is the normal motion flow (Alvarez *et al.*, 1993),

$$\frac{\partial I(x, y, \tau)}{\partial \tau} = \pm \|\nabla I\| \quad (7)$$

where the isophotes evolve along their normal direction with a constant speed and which corresponds to a morphological

erosion or dilation. Other examples include the so-called mean curvature flow (also called Euclidean shortening flow; Alvarez *et al.*, 1992),

$$\frac{\partial I(x, y, \tau)}{\partial \tau} = \|\nabla I\| \nabla \cdot \left( \frac{\nabla I}{\|\nabla I\|} \right), \quad (8)$$

in which the isophotes evolve along their normal direction with a speed proportional to their curvature, and the closely related affine shortening flow (Sapiro and Tannenbaum, 1993; Morel and Solimini, 1995)

$$\frac{\partial I(x, y, \tau)}{\partial \tau} = \|\nabla I\| \left[ \nabla \cdot \left( \frac{\nabla I}{\|\nabla I\|} \right) \right]^{1/3}. \quad (9)$$

Related diffusion schemes are the modified affine shortening flow (Niessen *et al.*, 1997) and entropy flow (Kimia and Siddiqi, 1996).

Geometry-based diffusion schemes can be extended to 3-D in which case the diffusion corresponds to an evolution of isophote surfaces. In the case of the mean curvature flow, the isophote surfaces evolve with a speed proportional to the mean curvature, i.e. the sum of the two principal curvatures. However, while in the two-dimensional (2-D) case any curve evolves to a round point without developing singularities, this is not the case for the 3-D counterpart. Moreover, geometry-based diffusion schemes cannot be applied to multi-valued images.

### 1.2.3. Anisotropic diffusion

Weickert (1997) has pointed out that the diffusion scheme proposed by Perona and Malik is governed by a scalar conductance term and the resulting flux is therefore isotropic. This has the disadvantage that the diffusion scheme tends to stop the diffusion at the edges rather than to permit diffusion along the edges. To allow for a truly anisotropic character of the diffusion equation, Weickert proposed a tensor anisotropic diffusion

$$\frac{\partial I(x, y, \tau)}{\partial \tau} = \nabla \cdot (\mathbf{C}(\nabla I) \nabla I) \quad (10)$$

where the eigenvectors  $\mathbf{e}_1$  and  $\mathbf{e}_2$  of the diffusion tensor  $\mathbf{C}$  are defined in such a way that  $\mathbf{e}_1 \parallel \nabla I$  and  $\mathbf{e}_2 \perp \nabla I$ . To encourage a diffusion along the edge rather than across the edge, Weickert suggested choosing the corresponding eigenvalues  $\lambda_1$  and  $\lambda_2$  as

$$\lambda_1 = g(\|\nabla I\|) \quad (11)$$

$$\lambda_2 = 1. \quad (12)$$

The resulting diffusion process is truly anisotropic and performs well in a wide variety of images (Weickert, 1998).

However, the equations do not take into account specific system characteristics such as inhomogeneous data sampling, and since the preferred directions for diffusion are dictated exclusively by the gradient of the image intensity, results are sensitive to image contrast and noise.

### 1.3. Overview

It has been pointed out that all of the above approaches rely only on the differential structure of the data, while ignoring the particular properties of the system (Sanchez-Ortiz *et al.*, 1996b; Sanchez-Ortiz, 1999). As a consequence, results still tend to be poor in regions of low contrast and low signal-to-noise ratio, where the differential characteristics of data do not provide enough information to distinguish between significant edges and noise. Some approaches oversimplify the conductance function of the diffusion process by ignoring the direction of the gradient or by disregarding important system characteristics such as non-homogeneous data sampling.

To overcome these problems we propose in this article a second-rank tensor conductance function with an explicit dependence on the space coordinates and the data function. This tensor modifies the nature of the diffusion equation making it heterogeneous and anisotropic from the starting point, and not as a consequence of the discretization scheme as in the case of some previous treatments. Furthermore, we develop a general framework to incorporate *a priori* knowledge of the geometry and dynamics of the system in vector-valued and multi-dimensional images.

We illustrate this approach with an application to cine volumetric anatomic and velocity-encoded MR images of the heart, where noise reduction in data and segmentation of the left-ventricle (LV) muscle (the myocardium) are commonly required.

A simplistic model for the LV shape and contraction rate is used to provide the equations with *a priori* information which is independent of noise. This information is combined with the differential characteristics of the data in order to produce a robust diffusion scheme. The diffusion process also combines information from vector- or multi-valued images, in this case anatomic and velocity-encoded data, thus helping the reconstruction of edges in regions where the anatomic image presents low contrast or signal-to-noise ratio.

Results of this and other diffusion schemes are compared on synthetic and MR images using quantitative evaluation tools like the correlation coefficient, and qualitative ones such as intensity profiles, visual inspection of reconstructed volumes and the GER-RGB device.

## 2. 4-D MULTI-VALUED MR DATA

Gated cine MR imaging techniques are synchronized with the electrocardiographic (ECG) signal and can be used to

generate a sequence of tomographic images of the heart where the image intensity is associated with the tissue type or ‘density’ ( $\rho$ ). The sequence of images correspond to different times ( $t$ ) which are usually distributed homogeneously during a cardiac cycle, covering most phases of the heartbeat. In addition to these sequences of 2-D images (with which we align the  $x$ - $y$  plane of the Cartesian coordinates reference system), multi-slice imaging can provide contiguous images parallel to the  $x$ - $y$  plane, at different ‘heights’ in a third spatial coordinate axis ( $z$ ). A common procedure for evaluating LV performance is to produce data such that  $x$  and  $y$  lie in the short-axis plane of the LV, and  $z$  along its long axis (see Figure 1 discussed in Subsubsection 3.3.1).

For each of the density images, velocity-encoded data in the same plane is acquired using a phase-sensitive MR technique. The velocity data is produced as three images,  $V_x$ ,  $V_y$  and  $V_z$ , which correspond to the Cartesian components of the velocity vector field  $\mathbf{V}$  all over the image domain.

We describe the data used as the vector function  $\mathbf{F}(\mathbf{p}) = (F_1(\mathbf{p}), \dots, F_n(\mathbf{p}), \dots, F_N(\mathbf{p}))$  where  $\mathbf{F} : \mathfrak{R}^M \rightarrow \mathfrak{R}^N$ ,  $F_n : \mathfrak{R}^M \rightarrow \mathfrak{R}$ ,  $\mathbf{p} = (p_1, \dots, p_m, \dots, p_M)$ , and  $p_m \in \mathfrak{R}$  (i.e.  $\mathbf{p} \in \mathfrak{R}^M$ ). In this example the space-time coordinates are  $x$ ,  $y$ ,  $z$ ,  $t$  ( $M = 4$ ), and the multiple features or values of the vector-valued image correspond to  $\rho$ ,  $V_x$ ,  $V_y$ ,  $V_z$  ( $N = 4$ ). Therefore, the vector data function takes the form

$$\mathbf{F}(\mathbf{p}) = (F_\rho(\mathbf{p}), F_{V_x}(\mathbf{p}), F_{V_y}(\mathbf{p}), F_{V_z}(\mathbf{p})) \quad (13)$$

where

$$\mathbf{p} = (x, y, z, t). \quad (14)$$

Within this framework we deal with  $t$  as a fourth ‘spatial’ coordinate and allow the cine sequence to provide extra information for the diffusion process. The method described below treats the space of coordinates as non-homogeneous and takes into consideration the different scales, physical units and sampling rates of the data in each of the coordinate axes.

The images employed in this work were acquired with a modified Picker 1.5 T MR scanner using electrocardiographic gating and a cine gradient-echo sequence. Eight contiguous image planes of  $256 \times 256$  pixels were used to cover the LV, each with a field of view of  $\sim 35$  cm and a slice thickness of 1 cm. The cine sequence consisted of 16 time frames distributed at intervals of 40 ms during the cardiac cycle. The diffusion process was applied to image windows of  $\sim 80 \times 80$  pixels in the short-axis plane ( $x$ - $y$ ), containing the entire LV.

The velocity data measurements were made on a range such that images contain meaningful information about blood and muscle tissue motion. Because blood velocity is much larger than muscle velocity in the direction of the long-axis

( $V_z$ ), standard wrap-around correction was used to adjust the value of blood velocity whose phase shift exceeded the measured velocity range (Underwood and Firmin, 1991). Since low-density material such as lung tissue produces meaningless random velocity values, all velocity components in the region of the lung were set to zero. This region was identified by thresholding the density image ( $\rho$ ), because in this type of image the lung tissue has a distinctive low-intensity value.

### 3. VECTOR-VALUED MULTI-DIMENSIONAL ANISOTROPIC DIFFUSION

In the case of this vector function of several variables that represents our four-dimensional (4-D) vector-valued data, the equation for anisotropic diffusion (AD) is

$$\frac{\partial \mathbf{F}(\mathbf{p})}{\partial \tau} = \nabla \cdot (\mathbf{C}^n(\mathbf{p}, \mathbf{F}) \nabla \mathbf{F}(\mathbf{p})), \quad (15)$$

i.e. the set of the four coupled equations:

$$\frac{\partial F_n(\mathbf{p})}{\partial \tau} = \nabla \cdot (\mathbf{C}^n(\mathbf{p}, \mathbf{F}) \nabla F_n(\mathbf{p})), \quad (16)$$

for all  $n \in \{\rho, V_x, V_y, V_z\}$ .

We must distinguish between the variables  $t$  and  $\tau$ . While the former is the time during the heartbeat and one of the coordinates in our 4-D set of data, the latter is the scale-space parameter that refers to the time during the diffusion process. The equations will be integrated with respect to  $\tau$ , from 0 to the total diffusion time  $\tau_{\text{total}}$ , and over all of the 4-D data domain. Following convention, the dependence of  $\mathbf{C}^n$  and  $\mathbf{F}$  on  $\tau$  has not been shown explicitly in the equations.

#### 3.1. Conductance function

The coupling term for these equations is the inhomogeneous and anisotropic conductance tensor function which we define by the product

$$\mathbf{C}^n(\mathbf{p}, \mathbf{F}) = l^n \mathbf{G}(\mathbf{F}) \mathbf{W}(\mathbf{p}), \quad (17)$$

where  $l^n$  are scalar constants used to control the diffusion rate for each of the feature images functions as described in Subsection 3.4.

The conductance function that we introduce here differs from previous definitions in two aspects. First, in contrast to other approaches (Nordstrom, 1990; Perona and Malik, 1990; Gerig *et al.*, 1992; Whitaker, 1993), this conductance function depends not only on the local behaviour of the data function  $\mathbf{F}$  which is subject to noise, but also on the particular characteristics of the system at every location  $\mathbf{p}$  of the 4-D coordinate space. The two factors that regulate the conductance—and therefore the diffusion—are, on the

one hand, the function matrix  $\mathbf{G}(\mathbf{F}(\mathbf{p}))$  that uses information obtained from the local behaviour of the data values; and, on the other hand, the weighting function matrix  $\mathbf{W}(\mathbf{p})$  obtained from *a priori* knowledge of the geometry and dynamics of the system from which the data is generated.  $\mathbf{W}$  is independent of the data function itself.

The second difference is that instead of a scalar function, the conductance function we define is a second-rank tensor. Such a definition permits real heterogeneous and anisotropic diffusion which is the result of the properties of the conductance function as well as of the data gradient of Equation (15). As we mentioned before, in the work of Perona and Malik the conductance was defined as a heterogeneous function over the image space, while the anisotropic character arose only as a consequence of their discretization scheme. In the present scheme, the tensors  $\mathbf{C}^n$ ,  $\mathbf{W}$  and  $\mathbf{G}$ , which expressed in matrix notation have entries of the form  $C_{ij}^n, W_{ij}, G_{ij} : \mathfrak{R}^4 \rightarrow \mathfrak{R}$  where  $i, j \in \{x, y, z, t\}$ , permit the conductance to vary not only according to location in the 4-D space but also depending on direction.

Although in this article we only examine the case in which the matrices are diagonal, a conductance function defined in this manner permits the diffusion process to be biased differently in every direction and in every position in space.  $\mathbf{G}$  accounts for the characteristics of the data. For instance, it could be a noise or boundary estimator.  $\mathbf{W}$  accounts for the intrinsic properties of the space. For example, a diffusion process may be modelled in which an electrical or gravitational field produces a force which varies with the direction—for an anisotropic system, and also varies with the distribution of mass or electrical charges in the material—for a heterogeneous medium.

### 3.2. Differential or data-based weights: $\mathbf{G}(\mathbf{F})$

The framework used to describe the function matrix  $\mathbf{G}$  is suitable for elaborated edge estimators. In this work, however, we use a 4-D vector-valued function analogous to the simple and widely used function  $g(\|\nabla I\|)$  of Equation (5). This monotonically decreasing scalar function of the magnitude of the gradient of  $I$  (the image intensity), has the desired effect of blurring small discontinuities, while sharpening edges when  $k$  is chosen adequately. The values of  $k_m$  (one for each coordinate axis) are computed using Canny's noise estimator (Canny, 1987): 90% of the value of the integral of the histogram of the gradient magnitude. The computation is made for every iteration of the diffusion process using the definition of the gradient magnitude shown below. Thus, the elements of the diagonal matrix  $\mathbf{G}$  are defined as follows:

$$G_{mm} = g_m(\|\nabla \mathbf{F}\|_m^*) = \frac{1}{1 + (\|\nabla \mathbf{F}\|_m^*/k_m)^2} \quad (18)$$

and

$$\begin{aligned} \|\nabla \mathbf{F}\|_m^* &\equiv \left( \sum_n (s_n J_{mn})^2 \right)^{1/2} \\ &= \left( \sum_n \left( s_n \frac{\partial F_n}{\partial p_m} \right)^2 \right)^{1/2}, \end{aligned} \quad (19)$$

for all  $m \in \{x, y, z, t\}$  and  $n \in \{\rho, V_x, V_y, V_z\}$ .

Since the data function that we use is a vector-valued function, we must define a norm for  $\mathbf{J}$ , the Jacobian matrix of  $\mathbf{F}$ , instead of using the magnitude of the gradient of  $\mathbf{F}$  as a dissimilarity measure. In this work we use the definition shown above as an alternative to the formerly proposed Euclidean norm of  $\mathbf{J}$  (Whitaker and Gerig, 1994). Such dissimilarity measure exploits the information provided by the different data feature functions  $F_n$ , while preventing homogeneous regions in some directions to shadow steep gradients in others. In the case of a single-valued function ( $N = 1$ ) the expression is reduced to that for the magnitude of the gradient of the image intensity.

The scaling constants  $s_n$  are used for standardizing the units of the different data feature functions. The constants are chosen to make dimensionless units and values within a similar range. Since the grey values of the velocity images have the same physical units ( $\text{cm s}^{-1}$ ), only the already dimensionless 'density' image has to be rescaled to match the range of the velocity values.

As we mentioned in Subsubsection 1.2.1, the standard anisotropic diffusion equation is known to present potential numerical instability. The conditions required for stability and for the existence and uniqueness of a solution have been discussed by several authors including Perona *et al.* (1994). In this work we use the functions  $g_m$  which are known to be stable, and a dissimilarity measure which is a special case of the Euclidean norm of  $\mathbf{J}$  (in fact the dissimilarity measure employed here treats each coordinate axis separately, in the same manner as standard anisotropic diffusion does). Since the Euclidean norm of  $\mathbf{J}$  is known to be well behaved, we can expect the matrix  $\mathbf{G}$  to be well posed and stable too. As for the influence of the knowledge-based weight matrix  $\mathbf{W}$ , as we shall see in Subsubsection 3.3, it is composed of smoothly varying multiplicative constants that only decrease the amount of diffusion and therefore are not a source of numerical instability. These expectations have been confirmed by our observations.

#### 3.2.1. Reference frame

In the case of the anisotropic diffusion scheme described in Subsubsection 1.2.3, the eigenvectors of the tensor

conductance are chosen accordingly with the direction of the local image gradient. Although these varying frames of reference seem to be a natural choice for the enhancement of edges, the correct estimation of these directions depends strongly on the amount of noise and edge contrast of the image. Also, since in this frame of reference the matrix is not diagonal, the method is expected to be computationally more expensive.

As for the diagonal matrix  $\mathbf{G}$  defined above, the reference frame is fixed and aligned with the coordinate axes along which the data was sampled. This is a suitable reference frame for many applications including that of the MR data used in this work (described in Section 2), because the data generation process relies heavily on it. For instance, the 4-D data is constructed as a collection of 2-D ( $x$ - $y$ ) images and the dimensions of a 4-D voxel (that correspond to a pixel) are normally different for every coordinate axis. Images belonging to different anatomical planes ( $z$ ) are usually created with a few minutes time interval during which the patient might move or the electromagnetic field can vary. Also, images corresponding to different times ( $t$ ) during the heart cycle are affected by non-homogeneous data sampling (see Subsection 3.3.2). Even within a 2-D image we can find inhomogeneities in the magnetic field or artefacts in the  $x$ - or  $y$ -axis due to the way in which tissue is encoded during data acquisition.

The present definition of  $\mathbf{G}$  (and  $\mathbf{W}$ ) deals with the image in its natural coordinate system and can easily incorporate simple or sophisticated models that take into account the anisotropic and inhomogeneous characteristics of the data, inherent to the imaging system. In this work we use the simple model introduced via the matrix  $\mathbf{W}$  as described in the following section, and we keep the functions  $g_m$  as close as possible to the well understood function used by Perona and Malik (the only differences are those necessary to incorporate vector-valued data such as the scaling constants  $s_n$ ). In the future we intend to investigate the combination of different reference frames. For instance, a variable frame for the matrix  $\mathbf{G}$  that follows the direction of the local gradient (as described in Subsubsection 1.2.3), and for the matrix  $\mathbf{W}$ , a fixed frame aligned with the coordinate axes of the image.

### 3.2.2. Multi-feature diffusion

The different types of feature images could be diffused in a completely independent manner, however previous work such as that of Gerig and Whitaker on  $T_1$ - and  $T_2$ -weighted MR images has shown that combining information from the different feature images can improve results. They have proposed to diffuse the feature images separately but using the conductance as a coupling term between the images, where the common image gradient is used to decide how

much to diffuse each image. Noise is filtered better because random fluctuations of the image intensity will rarely be registered in all feature images. Unlike noise, high gradient values corresponding to boundaries between different types of tissue tend to reinforce each other because they are likely to be present in all feature images. Using density and velocity data has similar advantages because blood flow differs from muscle motion and in some images, for instance in  $V_z$  during mid-systole or diastole, the boundaries between both tissue types are clearly marked and can help the density images to locate boundaries and prevent diffusion across them. On the other hand, density images contribute to locate boundaries in velocity images at times when motion is minimal, therefore preventing blood and muscle velocities from being diffused together.


While density and velocity images are different in nature, the velocity images are all components of the same velocity vector. As we just mentioned, we diffuse them separately but using the common conductance term which determines the amount of diffusion for all by estimating the presence of a boundary. This is a common procedure for diffusing vector images (Whitaker and Gerig, 1994) and in our case it is natural to deal with them as Cartesian components (instead of using magnitude and directions of the vectors, for instance), since this is the way in which the components are (separately) physically measured. An alternative way of clearing velocity data from noise would be to use fluid mechanics equations and information about the flow patterns, but this possibility is not explored in this work.

The next section describes how a weighting function for a specific problem can be chosen based on *a priori* knowledge of the system.

### 3.3. Knowledge-based weights: $\mathbf{W}(\mathbf{p})$

Information about the nature of the system and the data generation process can be very valuable when processing and analysing data. The system's geometry and dynamics, the sources of noise in the data, and in general, any kind of *a priori* knowledge can improve the results of tasks such as image segmentation or noise reduction. Using this information has a clear advantage over making statistical measurements of the regional properties of the data: it is noise independent. However, one must find a robust method to incorporate 'fresh' information from the data and thus avoid over-constraining the behaviour of the system and overlooking real characteristics which were not anticipated.

In particular, we are interested in all the information about the structure of the 4-D image that can be used to encourage intra- rather than inter-region diffusion. In the following we use simple knowledge about cardiac MR to generate a weight



**Figure 1.** Left, orientation of the coordinate system with respect to the barrel-shaped LV. Middle, cross-section of the LV and polar coordinates in the  $x$ - $y$  plane. Right,  $x$ - $z$  plane view of a pixel of the myocardium without a neighbouring pixel of the myocardium in the positive  $z$  direction.

function matrix  $\mathbf{W}(\mathbf{p})$  that reflects the characteristics of a given position within a, *per se*, heterogeneous and anisotropic space.

### 3.3.1. Geometric information

The LV has the approximate shape of a ‘capless’ ellipsoid and has been modelled using super-ellipsoids (Bardinet *et al.*, 1996) and other sophisticated deformable models (Park *et al.*, 1996), some of which used planispheric (Declerck *et al.*, 1997) or prolate coordinates (Hunter *et al.*, 1992; Matheny and Goldgof, 1995). In this article, however, we intend to introduce a minimal amount of *a priori* knowledge about the exact shape of the heart because one of the applications of the diffusion methods is precisely the segmentation and the consequent extraction of the detailed shape variations of the ventricle. Moreover, simple assumptions facilitate the portability of the method to other imaging environments that share similar characteristics.

A simple model that serves our purpose can be built from the realization that in the series of LV short-axis images described above, the myocardium exhibits a strong cylindrical symmetry. The left-hand diagram of Figure 1 shows the approximate shape of the middle section of the LV (excluding base and apex), and the orientation of a coordinate system that exploits this symmetry. When looking at the 2-D images in the  $x$ - $y$  plane, the symmetry is circular and the area of interest for the segmentation, the myocardium, appears as a ring formed by two almost concentric boundaries, as shown in the middle diagram of Figure 1. The outer boundary, the epicardium, has a shape that resembles a circle in all the slices and present little variation through the heart cycle. The shape of the inner boundary, or endocardium, varies

more with time. For approximately half of the images the shape is circular and for the other half the fluctuations could be regarded as moderate deformations of a circle. The most pronounced of the deformations occur in images where the papillary muscle<sup>a</sup> is indistinguishable from the endocardium.

Since MR provides high contrast between blood and tissue, the endocardium is normally preserved and emphasized with the part of the diffusion equation that uses the norm of the gradient of the intensity function. However, blood flow turbulence can considerably reduce the contrast and generate intensity variations that resemble tissue structures. In the case of the epicardium, the region contiguous to the lung is well defined, but the contrast between the heart and the rest of the (non-lung) tissue that surrounds it is poor and the diffusion process tends to merge them obliterating the boundary.

For these reasons we introduce the weighting factors  $W_{xx}$  and  $W_{yy}$  which acting together penalize diffusion within the  $x$ - $y$  plane in the radial direction of a polar coordinate system. The origin of the coordinate system is located at the centre of the circle that best fits the epicardium. In this manner we avoid blurring the outside boundaries of the myocardium and also the regions of the endocardium which are aligned with circles centred at the origin of such polar coordinate system (Figure 1, middle).

Finding the location of the centre of such a family of circles (one for each of the 2-D images) can be achieved by visual inspection, or can be done automatically with an algorithm based on the Hough (Ballard, 1981) transform or multi-scale medial analysis (Morse *et al.*, 1993). Since the

<sup>a</sup>Tubular-shaped muscles that are attached to the endocardium and cross part of the ventricular chamber to control the opening of the mitral valve.

method described is robust with respect to small errors in the location of this centre, we use the same coordinates of the centre for all the images,  $(x_{\text{origin}}, y_{\text{origin}})$ , making the weighting functions  $W_{xx}(x, y)$  and  $W_{yy}(x, y)$  independent of the coordinates  $z$  and  $t$ . The weighting factors are

$$W_{xx}(x, y) = |\sin(\theta(x, y))| \quad (20)$$

and

$$W_{yy}(x, y) = |\cos(\theta(x, y))| \quad (21)$$

where

$$\theta(x, y) = \arctan\left(\frac{y - y_{\text{origin}}}{x - x_{\text{origin}}}\right) \quad (22)$$

is the angle defined by the polar coordinate system in each of the slices.  $\theta$  also corresponds to the angle of the cylindrical polar coordinate system whose  $z$ -axis is the symmetry axis (the long axis) of the LV.

As for the third spatial coordinate axis ( $z$ ), we correlate pixels with a weight that depends on the average radius of the LV rings at every height on the  $z$ -axis. Contiguous planes where the radius of the myocardium ring is very different are less likely to have a  $z$ -neighbour that belongs to the myocardium (see right-hand diagram of Figure 1), and therefore should be less correlated. We find that the sinusoidal function

$$W_{zz}(z) = |\sin(\omega_z z + \phi_z)| \quad (23)$$

can approximately describe this correlation when adjusting the frequency ( $\omega_z$ ) and phase ( $\phi_z$ ) constants of the function according to the data (see the Appendix).

Although cylindrical coordinates would seem a natural choice for developing the equations (one could transform the data and Equation (15) to such coordinate system), the MR data is given in Cartesian coordinates and would require interpolation to convert it from one system to another. Since the interpolation could introduce unwanted blurring, we prefer to write and manipulate the equations directly in Cartesian coordinates.

### 3.3.2. Dynamic information

As for the fourth dimension, the time coordinate ( $t$ ), we can also correlate data with different weights using information about the behaviour of the heart during the heart cycle, specifically, using the fact that the LV contracts and expands at very different speeds during the various phases of systole and diastole. The immediate consequence of this is that data collected at times when the contraction rate is maximum is less correlated with data adjacent in the  $t$ -axis, i.e. the data gathered shortly before and afterwards in the cine sequence. On the other hand, data obtained when the heart

moves slowly should be strongly correlated in the time axis. This phenomenon can be viewed as non-homogeneous data sampling.

We can generate a function that describes precisely how to correlate data in the  $t$ -axis, either using textbook knowledge about the functional relationship between the phase and speed of the heart contraction and expansion, or using the velocity-encoded data to compute how fast specific regions of the myocardium move at different times of the sequence<sup>b</sup>. However, in this article we prefer to keep the model consistently simple and use the same type of functions in all coordinate axes. For this reason we overlook the fact that the heart cycle is not completely symmetrical between the contraction and expansion time, and approximate the correlation between successive data using the first harmonic of its Fourier series representation:

$$W_{tt}(t) = |\cos(\omega_t t + \phi_t)| \quad (24)$$

where the frequency ( $\omega_t$ ) and the phase ( $\phi_t$ ) are adjusted according to the data as described in the Appendix.

Equations (20)–(24) describe the elements of the diagonal function matrix  $\mathbf{W}(\mathbf{p})$  which incorporates knowledge from this simple model of the LV that has proved to be effective.

### 3.4. Total diffusion time for multi-feature data

As we mentioned before, there are two main reasons for using diffusion on images. If the purpose of the diffusion process is to pre-segment the images, a long diffusion time is suitable, as it will create homogeneous regions and sharp boundaries between them. However, when the purpose of the diffusion is to remove noise but without obliterating details of the local structure of the data, then the diffusion time must be small. In either case the diffusion time of the process ( $\tau_{\text{total}}$ ), which translates into the number of iterations in the discrete implementation, can be set interactively by visually inspecting the results or estimating its noise or data correlation. For example, Canny's noise estimator will gradually slow down the process by decreasing the value of  $k_m$  at an exponential rate (Simmons, 1992).

As discussed in Subsubsection 3.2.2, when using the multi-feature framework for diffusing vector-valued images, density and velocity data greatly benefit from each other. However, choosing a proper diffusion time is a delicate issue because of the conflicting motivation for diffusing each type of image. The most common reasons for diffusing density-encoded MR data are the enhancement and pre-segmentation of the images, therefore relatively long diffusion times are required. On the other hand, in the case of velocity-encoded

<sup>b</sup>See Morel and Solimini (1995) for a formal discussion of the enhancement of image sequences.



MR data a very short diffusion time is required because the main purpose of diffusing is to reduce some of the noise but without altering the regional variations of the velocity. These regional variations are of paramount importance to the analysis of non-rigid motion of the myocardium (Sanchez-Ortiz and Burger, 1995; Meyer *et al.*, 1996; Sanchez-Ortiz *et al.*, 1996a; Sanchez-Ortiz, 1999).

For these reasons we require a mechanism that controls the amount of diffusion which takes place in the density and velocity data when using the multi-feature framework. We found that the scalar constants  $l^n$  that multiply the conductance tensor in Equation (17) are a simple way to control the diffusion rate independently for each of the feature data functions  $F_n$ . In fact, in the discrete implementation these constants are also used to keep the scheme stable (see the Appendix).

In the next section we present the results of diffusing density- and velocity-encoded data, focusing particularly on relatively long diffusion times that make the process useful as a pre-segmentation tool.

#### 4. RESULTS

Equation (15) was solved using a 4-D finite-difference method analogous to the 2-D discretization scheme used by Perona *et al.* (1994; see the Appendix for details). The computational complexity of these algorithms is a linear function of the number of pixels of the  $m$ -dimensional image, and also of the number of iterations selected for the diffusion process. The weighting factors  $\mathbf{W}(\mathbf{p})$  were used to introduce adiabatic boundary conditions in the four coordinate axes, i.e. the conductance function was set to zero in the boundaries of the 4-D image<sup>c</sup>.

In this section we compare results of different types of diffusion applied to synthetic data and MR images of the LV. All images in a single figure were produced using the same parameters. This includes the total diffusion time<sup>d</sup> and, in the case of the anisotropic diffusion methods, the parameters  $k_m$  that were computed using Canny's noise estimator as described in Subsection 3.2. In the following, when referring to a particular image in a figure we use a suffix to distinguish each row and column, e.g. Figure 3. $r_2c_1$  refers to the first column and second row of Figure 3, and Figure 5. $c_4$  refers to the fourth column of Figure 5.

##### 4.1. 2-D synthetic data

In order to assess the performance of the proposed and other common diffusion schemes in cases where a ground truth is

known, we created simple synthetic images with a geometry that resembles that of the actual cardiac MR data and added to them a strong component of Gaussian noise.

Figures 2 and 3 show respectively the results of 25 and 75 iterations of diffusion for a synthetic image composed of concentric circles and added Gaussian noise. The first row shows, from left to right, the original image, the image after adding noise, the image processed with homogeneous isotropic diffusion (i.e. with a Gaussian filter), the image processed with the mean curvature flow (MCF) scheme described in Subsubsection 1.2.2, the image processed with 'standard' anisotropic diffusion (i.e. with the algorithm of Perona and Malik and the parameters  $k_m$  computed as described above), and the image processed with knowledge-based anisotropic diffusion. The second row shows the magnitude of the intensity gradient corresponding to each of the images on the row above.

Although all diffusion schemes smooth homogeneous regions of the image, Gaussian blur also affects the edges and makes their exact localization difficult (Perona *et al.*, 1994). The mean curvature flow and both anisotropic diffusion schemes tend to enhance boundaries and, as we can see in the gradient images, to locate them with reasonable precision and sharpness. However, the knowledge-based method has a better enhancing effect on the boundary of the inner circle, where the contrast was lower. This is not surprising since the scheme penalizes diffusion in the direction normal to the boundaries, and unlike the mean curvature flow method, the estimation of this direction does not depend only on the differential characteristics of the noisy data: it also relies on *a priori* information about the shape of the objects of interest in the image.


Figures 4 and 5 show respectively, the results of 25 and 75 iterations of diffusion for another synthetic image composed of squared regions and added Gaussian noise. Images are arranged in the same fashion as those in the figures with synthetic circular regions, and the algorithms and parameters used to create them are also the same as those in the previous figures. In particular, the *a priori* information used for the knowledge-based scheme is the same, i.e. a geometric model corresponding to circular shaped structures.

The purpose of the figure is to show the performance of all methods in sharp boundaries like the corners of the squares. The mean curvature flow scheme produces visually appealing smooth images with straight edges, however, the corners of the squares are strongly eroded<sup>e</sup>. On the other hand, the images produced with the anisotropic diffusion methods show


<sup>c</sup>Cine studies normally exclude late diastolic images. Otherwise one could take advantage of the cyclic nature of the heartbeat and treat the first and last cine images as contiguous in the time axis, in a kind of 4-D torus.

<sup>d</sup>The Gaussian filtered images were produced using the equivalent (linear) homogeneous isotropic diffusion scheme.

<sup>e</sup>As we will discuss below, the variation of the parameters that control the degree and speed of the smoothing was limited to the values where the scheme was not subject to numerical instability.



**Figure 2.** Synthetic circles ( $c_1$ ) with Gaussian noise ( $c_2$ ) and the results of 25 iterations of different diffusion schemes: isotropic diffusion ( $c_3$ ), mean curvature flow ( $c_4$ ), standard anisotropic diffusion ( $c_5$ ) and knowledge-based anisotropic diffusion ( $c_6$ ). The second row shows the gradients of the images on the first row.




**Figure 3.** Synthetic circles ( $c_1$ ) with Gaussian noise ( $c_2$ ) and the results of 75 iterations of different diffusion schemes: isotropic diffusion ( $c_3$ ), mean curvature flow ( $c_4$ ), standard anisotropic diffusion ( $c_5$ ) and knowledge-based anisotropic diffusion ( $c_6$ ). The second row shows the gradients of the images on the first row.

edges which are not so straight but corners that are better preserved.


We must observe that the performance of the knowledge-based method is basically the same as that of standard anisotropic diffusion, i.e. that the method respects characteristics of the image, such as boundaries in the corners of the squares, which do not conform with the simple geometric model designed for the retrieval of circular shapes. This highlights that the diffusion scheme is strongly governed by the data dependent non-linear term  $\mathbf{G}(\mathbf{F})$  of the conductance function, and that the knowledge-based term  $\mathbf{W}(\mathbf{p})$  does not

degrade the performance of the scheme by over-constraining the evolution of the image.

Another manifestation of the preference given to the tangential direction can be seen on the effect of diffusion on homogeneous regions with noise. When using the knowledge-based scheme, the randomly oriented patterns that we see in the gradient images become aligned with the circular regions. Although this is not a desired side effect, it is a very small drawback when compared with the good performance retrieving the boundaries we are interested in. On the other hand, our experience shows that



**Figure 4.** Synthetic squares ( $c_1$ ) with Gaussian noise ( $c_2$ ) and the results of 25 iterations of different diffusion schemes: isotropic diffusion ( $c_3$ ), mean curvature flow ( $c_4$ ), standard anisotropic diffusion ( $c_5$ ) and knowledge-based anisotropic diffusion ( $c_6$ ). The second row shows the gradients of the images on the first row.



**Figure 5.** Synthetic squares ( $c_1$ ) with Gaussian noise ( $c_2$ ) and the results of 75 iterations of different diffusion schemes: isotropic diffusion ( $c_3$ ), mean curvature flow ( $c_4$ ), standard anisotropic diffusion ( $c_5$ ) and knowledge-based anisotropic diffusion ( $c_6$ ). The second row shows the gradients of the images on the first row.

spurious image features are not formed more often with the proposed method than with standard anisotropic diffusion, and that its finite-difference implementation is subject to less numerical instability than that of the mean curvature flow method.

#### 4.2. SSD and CC

In order to quantitatively compare results obtained with the different diffusion schemes in the case when the ground truth is known, we have computed the squared sum of intensity differences (SSD) and the linear correlation coefficient (CC). For two images  $I_1(\mathbf{p})$  and  $I_2(\mathbf{p})$ , these quantities are defined

as

$$\text{SSD} = \frac{1}{n} \sqrt{\sum_{\mathbf{p} \in \mathcal{P}} (I_1(\mathbf{p}) - I_2(\mathbf{p}))^2} \quad (25)$$

and

$$\text{CC} = \frac{\sum_{\mathbf{p} \in \mathcal{P}} (I_1(\mathbf{p}) - \bar{I}_1(\mathbf{p}))(I_2(\mathbf{p}) - \bar{I}_2(\mathbf{p}))}{\sqrt{\sum_{\mathbf{p} \in \mathcal{P}} (I_1(\mathbf{p}) - \bar{I}_1(\mathbf{p}))^2} \sqrt{\sum_{\mathbf{p} \in \mathcal{P}} (I_2(\mathbf{p}) - \bar{I}_2(\mathbf{p}))^2}}, \quad (26)$$

where  $\bar{I}_1(\mathbf{p})$  and  $\bar{I}_2(\mathbf{p})$  denote the average intensities of the images over the image domain  $\mathcal{P}$ . The sum of square differences is zero only if both images are identical, and positive if the images differ (even if only by a scale factor). The

**Table 1.** Square sum of differences (SSD) and correlation coefficients (CC) obtained when comparing the gradient of the original circles image (Figure 2. $r_2c_1$ ) with all of the other gradient images of Figures 2. $r_2$  (for 25 iterations) and 3. $r_2$  (for 75 iterations).

	Number of iterations	Noised image	Isotropic diffusion	MCF diffusion	Standard AD	Knowledge-based AD
SSD	25	0.37	0.43	0.20	0.19	0.17
	75	0.37	0.64	0.22	0.16	0.14
CC	25	0.11	0.51	0.72	0.77	0.83
	75	0.11	0.38	0.68	0.81	0.86

**Table 2.** Square sum of differences (SSD) and correlation coefficients (CC) obtained when comparing the gradient of the original squares image (Figure 4. $r_2c_1$ ) with all of the other gradient images of Figures 4. $r_2$  (for 25 iterations) and 5. $r_2$  (for 75 iterations).

	Number of iterations	Noised image	Isotropic diffusion	MCF diffusion	Standard AD	Knowledge-based AD
SSD	25	0.41	0.44	0.13	0.14	0.14
	75	0.41	0.62	0.14	0.12	0.13
CC	25	0.16	0.56	0.87	0.86	0.86
	75	0.16	0.40	0.84	0.89	0.88

correlation coefficient can take values between zero and one: it is zero when the images are not correlated, and one when they are identical or a linear transformation of each other<sup>f</sup>.

Table 1 presents the SSD and CC values obtained when comparing the gradient of the original circles image (Figure 2. $r_2c_1$ ) with all of the other gradient images of Figures 2. $r_2$  and 3. $r_2$ . We can see that the correlation coefficient for the noised image is greatly increased with the diffusion methods and that among them, the knowledge-based scheme achieves the best results.

We can also appreciate that while both anisotropic schemes improve the correlation as the number of iterations increases, the isotropic and the MCF schemes decrease the correlation with longer diffusion times. It has been pointed out (Whitaker and Pizer, 1993) that due to the discrete nature of the finite-difference implementation, during the diffusion process edges undergo an initial period of gradient increase, followed by a stable stage, slow decay, and eventually, if the time of diffusion is too long, a period of rapid decay that leads to a uniform intensity image<sup>g</sup>. While the anisotropic diffusion schemes delayed the arrival of the rapid decay stage by means of the adaptive threshold parameters  $k_m$ , the mean curvature flow scheme lacks such a mechanism and reaches the decay stage sooner.

<sup>f</sup>In fact the CC can be negative and reach  $-1$  if the images have negative correlation, i.e. one is the ‘negative image’ of the other.


<sup>g</sup>Another possible scenario, the so-called ‘staircasing’ which arises when the values of the parameters  $k_m$  are too small, will be described in Subsection 4.3.

The moment when this decay stage is reached is also influenced by the parameters that control the degree and speed of the smoothing action of the mean curvature flow algorithm, namely, the kernel  $\sigma$  used for regularizing the gradient, and the time step used in the finite-difference implementation. The selection of those values was determined not only by the necessity to match the global performance of the other diffusion schemes<sup>h</sup>, but also by the necessity to avoid the range of values where the algorithm presented the numerical instability described in Niessen *et al.* (1997).

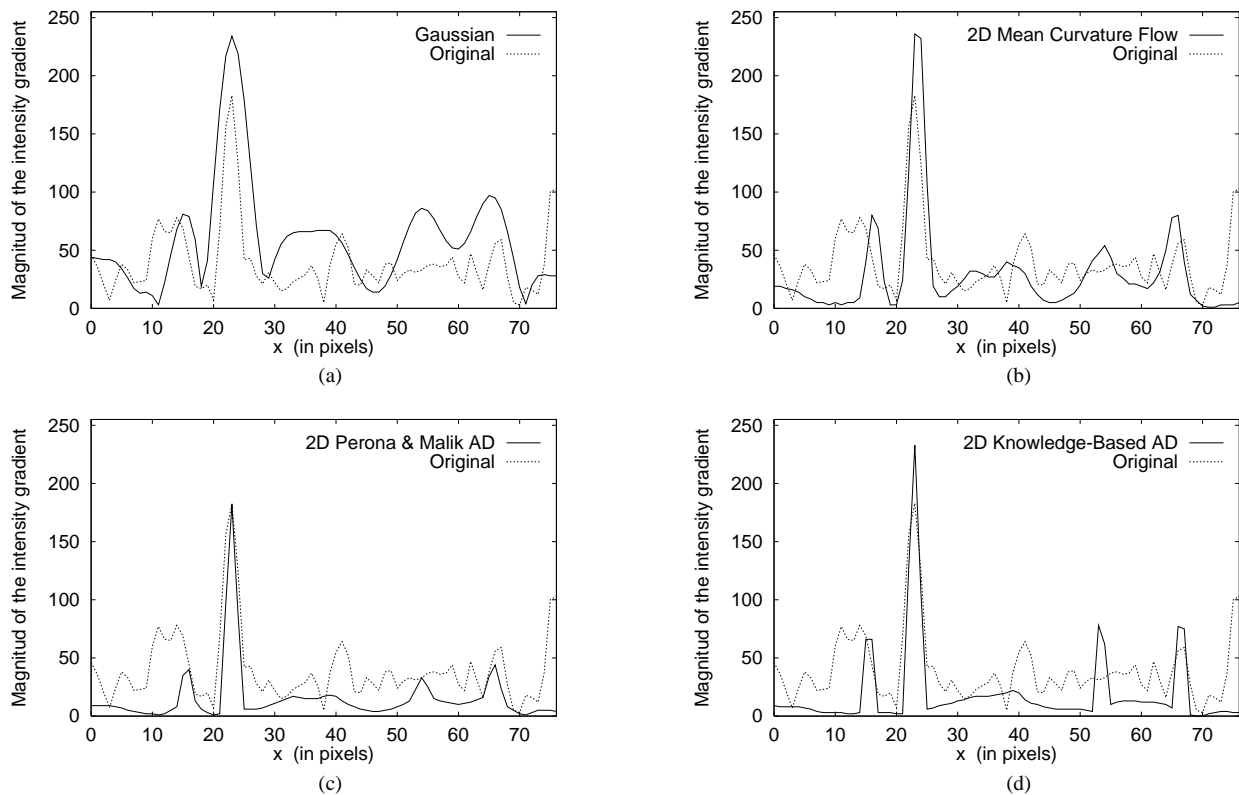
The SSD values shown in the tables provide a different measure that corroborates the results discussed for the correlation coefficients.

Table 2 presents the analogous of Table 1 for the square images of Figures 4 and 5. We can see that the SSD and CC values for the images processed with knowledge-based diffusion are very similar to those of the images processed with the standard anisotropic diffusion scheme. As we mentioned before, this shows that the proposed method performs reasonably well in situations where the image does not conform with the geometric model designed for the retrieval of circular shapes. We also notice that although mean curvature flow achieves slightly better values for short diffusion times, the performance is worse for a larger number of iterations.

<sup>h</sup>The degree of homogeneity reached in the background of the image assist in the assessment of the matching criterion.



**Figure 6.** Results of different diffusion schemes applied to an LV image. Top images, left to right: original, isotropic diffusion, mean curvature flow, standard anisotropic diffusion and knowledge-based anisotropic diffusion. Bottom: magnitude of the intensity gradient corresponding to each of the images in the row above. The horizontal lines in the middle of the images mark the regions for which intensity profiles have been plotted in Figure 7.



**Figure 7.** Profiles of the magnitude of the intensity gradient of the top images in Figure 6. The dashed line in all plots correspond to the intensity profile along the horizontal line in the middle of the original image gradient (Figure 6. $r_{2c_1}$ ). Solid lines are the intensity profiles along the horizontal line in the middle of the gradient images of Figure 6: (a) isotropic diffusion (Figure 6. $r_{2c_2}$ ); (b) mean curvature flow (Figure 6. $r_{2c_3}$ ); (c) standard anisotropic diffusion (Figure 6. $r_{2c_4}$ ); (d) knowledge-based anisotropic diffusion (Figure 6. $r_{2c_5}$ ).

### 4.3. 2-D MR data

Figure 6 presents the results of diffusing a 2-D systolic image of the myocardium. The first row shows, from left to right, the original image and the results of processing it with isotropic diffusion, mean curvature flow, standard anisotropic diffusion and knowledge-based anisotropic diffusion. The second row shows the magnitude of the intensity gradient corresponding to each of the images on the row above. The horizontal lines at the centre of the images mark the regions for which intensity profiles have been plotted in Figure 7.

Unlike Gaussian blur, mean curvature flow and standard anisotropic diffusion preserve and sharpen boundaries. However, in regions where the image intensity varies slowly, even these methods perform poorly. For instance, in Figure 6 we can see that when diffused with any of the first three methods, the right-hand side of the endocardium tends to disappear. It is only when processed with the knowledge-based technique that it is clearly outlined.

These differences can be better seen in Figure 7 where we compare the magnitude of the gradient of the intensity of the original and the diffused images along the horizontal line in the centre of the images of Figure 6. These plots show that the Gaussian filter crudely identifies the four edges along the line, although it also produces a protuberance in the region between pixels 30 and 40 where the blood movement makes the image intensity vary slowly.

Only the knowledge-based method smooths the homogeneous region and at the same time produces high and narrow peaks, i.e. a high-contrast image with sharply located edges. Mean curvature flow produces high but too wide peaks, making the exact localization of edges difficult. Although standard anisotropic diffusion smooths the homogeneous region and produces narrower peaks, these are not as high and therefore the image lacks contrast.

We must realize that this is not a consequence of having different ‘effective’ diffusion times in each of the diffusion methods. Every scheme, using the coefficients  $\lambda^n$  described in the Appendix, takes into account and compensates for the factors that influence the total flux of brightness during the diffusion process, namely, the number of pixel neighbours of the scheme (whether in 2-, 3- or 4-D), and the average values of the knowledge-based weights.

In order to illustrate the effect of using different diffusion time lengths, Figure 8 exhibits some stages of the process of diffusing a noisy early diastolic image. The first, third and fifth rows compare the results of mean curvature flow, standard and knowledge-based diffusion. The second, fourth and sixth rows show the enhanced magnitude of the gradients of the images on the rows immediately above. These gradients have been inverted and enhanced


for visual inspection by applying to them the function of Equation (5).

As we mentioned before, mean curvature flow lacks an adaptive mechanism like Canny’s noise estimator which helps the anisotropic diffusion schemes to enhance the contrast of significant boundaries, and therefore quickly reaches the stage where edges are not enhanced but eroded. On the other hand, anisotropic diffusion methods are liable to the undesired effect of ‘staircasing’ described by Whitaker and Pizer (1993), in which a single boundary is broken into many discrete steps. This occurs in slowly varying gradients when the value of  $k_m$  is too low and the diffusion time is too long. Although longer diffusion times help anisotropic diffusion to achieve sharper edges, in most of the figures presented in this article we use shorter periods (like that used to produce the images in the fourth column) and thus avoid deficiencies such as excessive blur and staircasing. Concerning the performance of the schemes in this figure, the main difference can be seen in low-contrast boundaries where the knowledge-based method reinforces the expected near circular shape of the regions.

### 4.4. GER-RGB

Since the original image and its corresponding gradient are noisy, it is difficult to assess from the intensity profile plots whether the edges have been accurately located. To enlighten this point and to further compare edge detection results we have developed the graphic device GER-RGB (graphic edge representation from red, green and blue) (Sanchez-Ortiz, 1999). This representation facilitates the comparison of the edges’ positions and sharpness between two or three similar images. The images could be, for instance, parts of an animated sequence, or as in the case of this article, a single image which has been processed by different methods. The magnitude of the gradient intensity of each of the three images is thresholded and three binary images are produced. The threshold can be set by inspection or automatically selected using Canny’s noise estimator. The binary images are coloured with red, green and blue respectively, and combined to produce an RGB colour map that can be overlaid on one of the original images. As we shall see in the next figure, the colours in the new image provide ‘labels’ for each pixel that indicate whether the pixel was selected by the threshold in one or more of the red, green and/or blue images.


Figure 9 employs the above-mentioned device to compare the results obtained with the different diffusion schemes. The first and second rows show, respectively, the images of Figure 6 and their corresponding enhanced gradients. Here we observe that knowledge-based diffusion provides better defined edges and less areas where edges have not been properly detected nor properly erased either. In the




**Figure 8.** Different stages of the diffusion process. The columns, from left to right, show the original and the diffused images after 12, 25, 50, 100 and 200 iterations respectively. The first, third and fifth rows compare the results of mean curvature flow, standard and knowledge-based anisotropic diffusion respectively. The second, fourth and sixth rows show the enhanced magnitude of the gradients of the images on the rows immediately above.

third row we compare the edge detection performance of the schemes using the GER-RGB device. The magnitude of the intensity gradient of each diffused image (Figure 6.r<sub>2</sub>) has

been thresholded and the resulting binary images have been coloured and overlaid onto the original MR image. Thus, the second column presents results of mean curvature flow



**Figure 9.** Comparison of images diffused with mean curvature flow, standard anisotropic diffusion and 2-D knowledge-based anisotropic diffusion, using the GER-RGB device. The last row shows the coloured thresholded gradients for each diffused image, and its RGB colour superposition on columns 5 and 6.



**Figure 10.** Comparison of images diffused with mean curvature flow, standard and 2-D knowledge-based anisotropic diffusion, using the GER-RGB device. The last row shows the coloured thresholded gradients for each diffused image, and its RGB colour superposition on columns 5 and 6.



(image in blue), the third column of standard anisotropic diffusion (red) and the fourth column of the knowledge-based scheme (green).

The fifth and sixth columns show respectively the red–green and red–green–blue superpositions of the colour thresholds described above. Figure 9.*r3c5* compares and illustrates the difference in performance of standard and knowledge-based anisotropic diffusion. The yellow areas correspond to the regions that both methods have preserved as edges. The red and green areas are those that only one of the methods has preserved (red for standard and green for knowledge-based diffusion). It is clear that knowledge-based diffusion preserves many regions of the boundary of the myocardium that are eroded by standard diffusion. Image 9.*r3c6* adds the blue component corresponding to the mean curvature flow image, changing the colour of the commonly detected areas as described by the palette in the figure.

Figure 10 presents results for another LV image in a plane near to the apex. The images were produced and placed in the same fashion as the previous figure. The second, third and fourth columns show results for mean curvature flow, standard and knowledge-based diffusion respectively. It can be clearly seen that the latter preserves and emphasizes the left-hand side of the epicardium better than the other schemes. The colour images in the fifth and sixth columns compare the exact location of the enhanced edges.

In Figure 10.*r3c6* white regions correspond to significant edges marked in all of the coloured images. The blue areas surrounding white edges show that MCF does not locate edges as sharply as anisotropic diffusion methods do. Green edges are those marked only by the proposed method. They are formed as a consequence of the increased flux of brightness along the preferred directions of the knowledge-based model, i.e. mostly along the perimeter of the concentric circles described in Subsubsection 3.3.1. The plateau of intensity that corresponds to the myocardium, being aligned with such circles, is gradually reinforced by this flux of brightness and therefore discriminated from other regions in the image.


In the case of the top-right section of the epicardium, however, the algorithms are not so successful. This is most likely due to the pronounced difference of intensities in a large region of the myocardium in this area. Although human vision can easily compensate such intensity variations and discern an apparently continuous epicardium, attentive inspection reveals a dark region of the LV near to the uppermost part of the lung (the lung appears as the darkest shade of grey in all the images). The effect of this region is to ‘interrupt’ the flux of brightness along the myocardium and hinder the building of the edge. We will see another example of this effect in Figure 12.

Mean curvature flow produces visually appealing smooth and continuous images, however, edges are also blurred unless the algorithm is used only with few iterations. On the other hand, the implementation of Perona and Malik’s algorithm with the adaptive threshold based on Canny’s noise estimator provides sharper edges and can be used for longer diffusion times without blurring these edges. Since the parameters used by both of the anisotropic diffusion schemes are mostly the same, and because the characteristics of the images produced by them are similar, it is more meaningful to compare results of the proposed scheme to those obtained with standard anisotropic diffusion. For these reasons in the following figures we omit results of the mean curvature flow algorithm and focus on the results of the anisotropic diffusion schemes, thus highlighting the differences introduced by the knowledge-based method proposed.

#### 4.5. Image sequences

In Figure 11 we analyse the particularly noisy image used in Figure 8. The images are arranged similarly to those of the previous figures. The first four columns correspond, respectively, to the original image, the image diffused with 2-D standard anisotropic diffusion, with 2-D knowledge-based anisotropic diffusion, and finally, with 3-D (2-D+ $T$ ) knowledge-based anisotropic diffusion. The latter was produced by diffusing together four consecutive images of the cine sequence. Every diffused image in the figure was produced using 50 iterations. Images in the fifth and sixth columns compare the edge-detection performance of these methods using the GER-RGB device. The colours blue and red have been assigned to the edges of the original image and the standard diffusion results respectively. The colour green was used for 2-D knowledge-based diffusion in the third row and for the 3-D homologous in the fourth row.

As shown in Figure 11.*r3c3*, the knowledge-based method produce a nearly continuous boundary of the myocardium. The fourth column shows that 3-D data can further improve the results of anisotropic diffusion. We also notice that Figure 11.*r4c4* shows a double edge for the bottom left-hand region of the endocardium, one includes and the other excludes the papillary muscle. This is a result of using information belonging to consecutive images in which the myocardium has already relaxed after the contraction and there is a small gap between the papillary muscle and the endocardium (see Figure 12.*c1*). Plausible explanations for the appearance of the gap in the data sequence are either that muscle dilation brought them apart, or that due to the through-plane movement of the heart, the images correspond to a slightly different section of the ventricle (Sanchez-Ortiz and Burger, 1995). In either case, effects such as double-edge enhancing, which are introduced when using




**Figure 11.** GER-RGB comparison between images diffused with standard anisotropic diffusion (AD), 2-D knowledge-based (KB) AD, and 3-D KB AD. The third row shows the coloured thresholded gradients for the original, standard and 2-D KB AD image, and its RGB colour superposition on columns 5 and 6. The fourth row shows analogous results but using the 3-D KB AD image instead of the 2-D KB AD image shown in the row above. A colour version of this figure can be found on the CD-ROM.

multi-dimensional data, could be minimized using more detailed knowledge of the system dynamics.

In the colour images of Figure 11.c<sub>5</sub> we compare the results of the knowledge-based method to those of standard diffusion. In Figure 11.c<sub>6</sub>, adding the blue component coming from the gradient of the original image demonstrates that both methods locate the boundaries accurately (white lines show overlapping edges), and also that using *a priori* knowledge can help retrieving edges with gradient values lower than the average noise in the image (as shown by the pure green lines, where an edge has been detected and there is no ‘blue trace’ of it).

Figure 12 shows the sequence of four images (columns 1 and 2) used to produce the 3-D diffused images of Figure 11, as well as the results of applying to them the standard 2-D anisotropic diffusion (columns 3 and 4) and the 3-D (2-D+T) knowledge-based scheme (columns 5 and 6). The

images in the second row are those previously shown in Figure 11. The results of both methods are similar when the quality of the original image is good, for instance for the endocardium and the ventricle chamber in the last two rows. However, when the image quality is poor, as in the case of the endocardium and ventricle chamber of the first two rows, the proposed scheme performs much better. This is also the case for the bottom region of the epicardium in the last two rows. Figure 12.r<sub>3c1</sub> shows particularly low intensity values of the bottom region of the myocardium, which result in low contrast for the epicardium. There, we can also see a small ‘intensity notch’ in the epicardium. The effect of the notch is to diminish the flux of brightness along the myocardium and hinder the enhancement of the edge, in a manner similar to that described previously when discussing Figure 10. However, although the low contrast and the notch in the edge severely affect the performance



**Figure 12.** Results on a sequence of density MR images of the myocardium. Each row corresponds to a different time during the sequence. The original images are shown in the first column. The third and fifth columns show the images diffused with 2-D Perona and Malik, and 3-D knowledge-based anisotropic diffusion respectively. Even columns show the enhanced magnitude of the intensity gradients of the images on the previous column.

of standard diffusion (see Figures 12. $r_3c_3$  and 12. $r_3c_4$ ), 3-D knowledge-based diffusion enhances the edge as expected (see Figures 12. $r_3c_5$  and 12. $r_3c_6$ ).

#### 4.6. Volumetric data


Figure 13 shows the isosurface volume rendering of the LV endocardium produced from density-encoded data after five (left) and 50 (right) iterations of the proposed diffusion process. The data used for this figure has a long-axis resolution somewhat higher than that used in the other figures (14 planes instead of eight).

Figure 14 shows the isosurface volume rendering of a diffused cine sequence of the LV endocardium. In this case 8 planes and 16 time frames were used. We believe that the results are encouraging for a data set with low resolution

in the vertical axis and for the very simple method used for reconstructing the volumes. Section 6 describes the animated sequences shown in the CD-ROM made from this volumetric cine sequence.

#### 4.7. Velocity data

The evaluation of diffusion results for multi-feature data is a complex subject. As we mentioned in Subsection 3.4, velocity data is normally diffused for short periods in order to preserve information about the local non-rigid motion of the myocardium; therefore sophisticated methods are necessary to assess its efficiency (Meyer *et al.*, 1996; Sanchez-Ortiz *et al.*, 1996a). For this reason we present only qualitative results that highlight the advantages of using our scheme for noise reduction, and leave the thorough analysis of velocity



**Figure 13.** Volume rendering of the LV endocardium produced from density-encoded data after five (left) and 50 (right) iterations of the proposed diffusion process.

data for a forthcoming article in which we also make more extensive use of the dynamic information provided by the velocity vector field.

Figure 15 shows the results of knowledge-based anisotropic diffusion applied to a sequence of density- and velocity-encoded data. Even rows (2, 4 and 6) show the diffused version of the original images shown in the rows above (1, 3 and 5 respectively). Time frames of the sequence are arranged in normal English reading order. They begin with early-systole (top rows), continue with end-systole (middle rows) and finish with diastole (bottom rows). The same values of the diffusion rate constants  $l^n$  have been used for all feature images; hence, the diffused density images give good indication of the short diffusion time used on the velocity data (10 iterations).

Arrows show the direction of the in-plane components of the velocity vectors at some points of the myocardium. Few points were used in order to avoid cluttering the image with arrows. The points selected for sampling the velocity field are homogeneously distributed on the perimeter of an ellipse. The (nearly circular) ellipse was arbitrarily selected in an interactive manner in order to guarantee that it was fully contained in the region of the myocardium. For a given time frame, the same points of the myocardium were selected on the original and diffused data.

We can see that the velocity vectors in the diffused images show much more coherence than those in the original images. Randomly directed arrows have been eliminated, however the enhanced velocity vectors are not completely aligned with


each other and thus provide information about the non-rigid motion of the cardiac muscle.

## 5. CONCLUSIONS

Two fundamental aspects of this model are introduced by defining the conductance function as a second-rank tensor and as an explicit function of position and time. First, an explicit function of position and time can treat the space as intrinsically heterogeneous and anisotropic, while incorporating available *a priori* knowledge about the system. Secondly, defining the conductance as a second-rank tensor allows the model to bias different directions in the coordinate space, making the diffusion truly anisotropic and permitting the use of known symmetries of the system to improve the identification of the image boundaries.

Most of the *a priori* knowledge that we have incorporated in this application relies on the simple realization that the LV has nearly cylindrical symmetry. We have used weighting functions that encourage diffusion in the directions where it is most likely to find pixels of the same tissue type, i.e. along concentric circles in the 2-D images. When diffusing image sequences the weights have been adjusted to compensate for inhomogeneous data sampling using knowledge about the rates of contraction during the different phases of the cardiac cycle.

The model could be made more realistic either by using an ellipsoid, super-ellipsoid (Bardinet *et al.*, 1996), or incorporating detailed information about the ongoing movement and




**Figure 14.** Volume rendering of a cine sequence of the LV endocardium after 80 iterations of the proposed diffusion process.

deformations of the heart as assessed by non-rigid motion analysis methods (Sanchez-Ortiz and Burger, 1995; Park *et al.*, 1996; Sanchez-Ortiz *et al.*, 1996a). For instance, we are currently studying the effect of using velocity data to track regions of the heart in a cine sequence in order to correlate them during the diffusion process. Also, the knowledge-

based weights could explicitly depend on the type of feature data in order to exploit the specific characteristics of each type of images.

Although Canny's noise estimator is a simple and effective way to set a value for the parameters  $k_m$ , a histogram computed over the whole image misses gradients below the



**Figure 15.** Results of diffusion on a cine sequence of density- and velocity-encoded images. Odd and even rows show the original and diffused data respectively.

critical value and therefore hinders the detection of faint edges. Knowledge-based diffusion improves the detection of such edges, however, a local measure of the noise would further improve the performance of the scheme.

The subject is rich in promising possibilities and we are at present expanding this work in several directions. Among them are the use of velocities to extract information about the complex motion of the heart, and the use of adaptive weights produced by fuzzy clustering coupled with the diffusion process (Sanchez-Ortiz, 1998). Also, preliminary results show that the scheme can be successfully adapted to other imaging environments like ultrasound images of the prostate.

## 6. CONTENTS OF CD-ROM VIDEO

Contents of the video sequence:

- (i) **Cine volumetric data before diffusion.** Animation of the isosurface volume rendering of the cine volumetric data discussed in Subsection 4.6, after 10 iterations of the diffusion process.
- (ii) **Cine volumetric data after diffusion.** Animation of the isosurface volume rendering of the cine volumetric images shown in Figure 14, Section 4.6, after 80 iterations of the diffusion process.
- (iii) **Volumetric data during diffusion.** Animation of the isosurface volume rendering of early systole volumetric data during the diffusion process.
- (iv) **Surface rendering of cine volumetric data.** We show the surface rendering of a cine sequence of the LV endocardium. The data was first pre-segmented using the proposed diffusion scheme, then automatically segmented using a 3-D tracking algorithm (Rueckert and Burger, 1997) based on deformable models, and finally rendered as a sequence of surfaces. The value of the parameter that controls the rigidity of the deformable model was set high in order to include the papillary muscle within the LV chamber. The abrupt changes seen in some regions of a surface are due to the low resolution of the data in the vertical axis (eight planes). We believe the results are encouraging for an automated method in which the only parameters to control are the diffusion time and the rigidity of the deformable model.

## ACKNOWLEDGEMENTS

We thank many members of the MR Unit of the National Heart and Lung Institute of the Royal Brompton Hospital, London, for providing MR data. GIS-O acknowledges the

grant given by CONACYT. We also thank the referees for their very useful suggestions.

## REFERENCES

- Alvarez, L., Lions, P. L. and Morel, J. M. (1992) Image selective smoothing and edge detection by nonlinear diffusion II. *SIAM J. Numer. Anal.*, 29, 845–866.
- Alvarez, L., Guichard, F., Lions, P. L. and Morel, J. M. (1993) Axioms and fundamental equations of image processing. *Archive Rational Mech. Anal.*, 123, 199–257.
- Ballard, D. H. (1981) Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recogn.*, 13, 111–122.
- Bardinet, E., Cohen, L. D. and Ayache, N. (1996) Tracking and motion analysis of the left ventricle with deformable superquadrics. *Med. Image Anal.*, 1, 129–149.
- Canny, J. (1987) A computational approach to edge detection. *IEEE Trans. PAMI*, 8, 679–698.
- Catté, F., Lions, P. L., Morel, J. M. and Coll, T. (1992) Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.*, 29, 182–193.
- Declerck, J., Feldmar, J. and Ayache, N. (1997) Definition of a 4D continuous polar transformation for the tracking and the analysis of LV motion. In *Proc. Int. Conf. on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'97)*, Grenoble, Vol. 1205, pp. 33–42.
- Gerig, G., Kübler, O., Kikinis, R. and Jolesz, F. A. (1992) Nonlinear anisotropic filtering of MRI data. *IEEE Trans. Med. Imag.*, 11, 221–231.
- Hunter, P. J., Nielsen, P. M. F., Smaill, B. H., LeGrice, I. J. and Hunter, I. W. (1992) An anatomical heart model with applications to myocardial activation and ventricular mechanics. *Critical Rev. Biomed. Eng.*, 20, 403.
- Kimia, B. B. and Siddiqi, K. (1996) Geometric heat equation and nonlinear diffusion of shapes and images. *Comp. Vision Image Understanding*, 64, 305–322.
- Koenderink, J. J. (1984) The structure of images. *Biol. Cybernet.*, 50, 363–370.
- Lindeberg, T. and ter Haar Romeny, B. M. (1994) Linear scale space I: basic theory. In ter Haar Romeny, B. M. (ed.), *Geometry-Driven Diffusion in Computer Vision*, Ch. 1, pp. 1–38. Kluwer Academic Publishers, Dordrecht.
- Matheny, A. and Goldgof, D. B. (1995) The use of 3-dimensional and 4-dimensional surface harmonics for rigid and nonrigid shape recovery and representation. *IEEE Trans. PAMI*, 17, 967–981.
- Meyer, F. G., Constable, R. T., Sinusas, A. J. and Duncan, J. S. (1996) Tracking myocardial deformation using phase contrast MR velocity fields: a stochastic approach. *IEEE Trans. Med. Imag.*, 15, 453–465.
- Morel, J. M. and Solimini, S. (1995) *Variational Methods for Image Segmentation*. Birkhäuser, Boston, MA.

- Morse, B. S., Pizer, S. M. and Liu, A. (1993) Multiscale medial analysis of medical images. In Barret, H. H. and Gmitro, A. F. (eds), *Information Processing in Medical Imaging, Lecture Notes in Computer Science*, Vol. 687, pp. 112–131. Springer-Verlag, Berlin.
- Niessen, W. J., ter Haar Romeny, B. M., Florack, L. M. J. and Viergever, M. A. (1997) A general framework for geometry-driven diffusion equations. *Int. J. Comp. Vision*, 21, 187–205.
- Nordstrom, K. N. (1990) Biased anisotropic diffusion: a unified regularization and diffusion approach to edge detection. *Image Vision Comput.*, 8, 318–327.
- Park, J., Metaxas, D. and Axel, L. (1996) Analysis of left ventricular wall motion based on volumetric deformable models and MRI-SPAMM. *Med. Image Anal.*, 1, 53–71.
- Perona, P. and Malik, J. (1990) Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. PAMI*, 12, 629–639.
- Perona, P., Shiota, T. and Malik, J. (1994) Anisotropic diffusion. In ter Haar Romeny, B. M. (ed.), *Geometry-Driven Diffusion in Computer Vision*, Ch. 3, pp. 73–92. Kluwer Academic Publishers, Dordrecht.
- Rueckert, D. and Burger, P. (1997) Shape-based segmentation and tracking in 4D cardiac MR images. In *Proc. Int. Conf. on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'97)*, Grenoble, pp. 43–52.
- Sanchez-Ortiz, G. I. (1998) Fuzzy clustering driven anisotropic diffusion: enhancement and segmentation of cardiac MR images. In *IEEE Nuclear Science Symp. and Medical Imaging Conf.*, Toronto.
- Sanchez-Ortiz, G. I. (1999) *Knowledge-based Anisotropic Diffusion and Non-rigid Motion Analysis of Anatomic and Velocity Encoded Cardiac MR Images*. Ph.D. Thesis, Imperial College, London. To appear.
- Sanchez-Ortiz, G. I. and Burger, P. (1995) Vector field analysis of the dynamics of the heart using velocity encoded NMR images. In *9th Int. Symp. on Computer Assisted Radiology, CAR'95*, Berlin, pp. 228–233. Springer-Verlag, Berlin.
- Sanchez-Ortiz, G. I., Rueckert, D. and Burger, P. (1996a) Motion and deformation analysis of the heart using thin-plate splines and density and velocity encoded MR images. In *16th Leeds Ann. Statistical Research Workshop on Image Fusion and Shape Variability Techniques, LASR'96*, Leeds, pp. 71–78.
- Sanchez-Ortiz, G. I., Rueckert, D. and Burger, P. (1996b) Knowledge-based anisotropic diffusion of vector-valued 4-dimensional cardiac MR images. In *British Machine Vision Conf., BMVC'96*, Edinburgh, pp. 605–614.
- Sapiro, Z. and Tannenbaum, A. (1993) Affine invariant scale-space. *Int. J. Comp. Vision*, 11, 25–44.
- Simmons, A. (1992) *Segmentation of Neuroanatomy in Magnetic Resonance Images*. Ph.D. Thesis, Department of Medical Physics and Bioengineering, University College London, University of London.
- ter Haar Romeny, B. M. (ed.) (1994) *Geometry-driven diffusion in computer vision. Computational Imaging and Vision*. Kluwer Academic Publishers, Dordrecht.
- Underwood, R. and Firmin, D. (1991) *Magnetic Resonance of the Cardiovascular System*, 1st edn. Blackwell Scientific Publications, Oxford.
- Weickert, J. (1997) A review of nonlinear diffusion filtering. In *Scale-Space Theory in Computer Vision*, Utrecht, The Netherlands, pp. 3–28.
- Weickert, J. (1998) *Anisotropic Diffusion in Image Processing*. Teubner-Verlag.
- Whitaker, R. T. and Gerig, G. (1994) Vector-valued diffusion. In ter Haar Romeny, B. M. (ed.), *Geometry-Driven Diffusion in Computer Vision*, Ch. 4, pp. 93–134. Kluwer Academic Publishers, Dordrecht.
- Whitaker, R. T. and Pizer, S. M. (1993) A multi-scale approach to nonuniform diffusion. *Comp. Vision Graphics Image Process.*, 57, 99–110.
- Whitaker, R. T. (1993) Geometry-limited diffusion in the characterization of geometric patches in images. *Comp. Vision Graphics Image Process.*, 57, 111–120.
- Witkin, A. (1983) Scale-space filtering. In *Int. Joint Conf. on Artificial Intelligence*, pp. 1019–1022.

## APPENDIX A. DISCRETIZATION SCHEME

Equation (15) can be discretized on a hyper-cubic lattice (the 4-D analogous of a 2-D square lattice), where the data value is associated with the vertices, and conduction coefficients with the arcs. Using a 4-D extension of the four-nearest-neighbours discretization scheme employed by Perona *et al.* (1994), we obtain the following scheme:

$$\mathbf{F}^\tau(\mathbf{p}) \Big|_{\mathbf{p}=\mathbf{p}_0}^{\tau=i+1} = \mathbf{F}^\tau(\mathbf{p}) \Big|_{\mathbf{p}=\mathbf{p}_0}^{\tau=i} + \lambda^n \left( \sum_{d \in \mathcal{D}} C^d(\mathbf{p}, \mathbf{F}) \Delta^d \mathbf{F}^\tau(\mathbf{p}) \right) \Big|_{\mathbf{p}=\mathbf{p}_0}^{\tau=i}. \quad (\text{A1})$$

These are in fact four coupled equations since  $\mathbf{F} \rightarrow F_n$ , where  $n \in \{\rho, V_x, V_y, V_z\}$ . The computations are carried out for every point  $\mathbf{p}_0 = (x_0, y_0, z_0, t_0)$  of the data domain. The new value of  $\mathbf{F}$  (the diffused data) is calculated for the next iteration step ( $\tau = i + 1$ ) from the data at the present stage ( $\tau = i$ ).

Since the matrix  $\mathbf{C}$  is diagonal, we simplify the nomenclature by introducing the direction index  $d$ . Then, the summation takes place over all the eight neighbouring directions  $\mathcal{D} = \{\text{E, W, N, S, U, D, A, B}\}$  where the letters stand for East and West (E, W) on the  $x$ -axis, North and South (N, S) on the  $y$ -axis, up and down (U, D) on the  $z$ -axis, and after and before (A, B) on the  $t$ -axis. The index  $n$  that showed the dependence of  $\mathbf{C}$  on the type of feature data has been directly attached to the coefficients  $\lambda^n$  that control the intensity flux described for Equation (A13).



The following definitions hold for the present scheme:

$$(\Delta^d \mathbf{F}^\tau(\mathbf{p})) \Big|_{\mathbf{p}=\mathbf{p}_0}^{\tau=i} = \mathbf{F}^\tau(\mathbf{p}) \Big|_{\mathbf{p}=\mathbf{p}_0+\mathbf{d}\mathbf{p}^d}^{\tau=i} - \mathbf{F}^\tau(\mathbf{p}) \Big|_{\mathbf{p}=\mathbf{p}_0}^{\tau=i} \quad (\text{A2})$$

where

$$\mathbf{d}\mathbf{p}^d = (\delta^{Ed} - \delta^{Wd}, \delta^{Nd} - \delta^{Sd}, \delta^{Ud} - \delta^{Dd}, \delta^{Ad} - \delta^{Bd}) \quad (\text{A3})$$

and  $\delta$  is the Kronecker delta defined as

$$\delta^{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j. \end{cases} \quad (\text{A4})$$

The conductance function is computed as

$$C^d(\mathbf{p}, \mathbf{F}^\tau) = W^d(\mathbf{p}) g_d(\|\nabla \mathbf{F}^\tau(\mathbf{p})\|_d^*) \quad (\text{A5})$$

where  $g_d$  is that of Equation (18), and Equation (19) becomes

$$\|\nabla \mathbf{F}^\tau(\mathbf{p})\|_d^* = \left( \sum_n (s_n \Delta^d F_n^\tau(\mathbf{p}))^2 \right)^{1/2}, \quad (\text{A6})$$

for all  $d \in \mathcal{D}$  and  $n \in \{\rho, V_x, V_y, V_z\}$ .  $\Delta^d$  is defined for the single-valued functions  $F_n$  in a fashion similar to that of its vector homologous of Equation (A2). In this notation, the weights functions can be expressed as

$$W^d(\mathbf{p}) = |W_{xx}^d(x, y) + W_{yy}^d(x, y) + W_{zz}^d(z) + W_{tt}^d(t)|, \quad (\text{A7})$$

where

$$W_{xx}^d(x, y) = (\delta^{Ed} + \delta^{Wd}) \times \sin\left(\theta\left(x + \frac{1}{2}(\delta^{Ed} - \delta^{Wd}), y\right)\right), \quad (\text{A8})$$

$$W_{yy}^d(x, y) = (\delta^{Nd} + \delta^{Sd}) \times \cos\left(\theta\left(x, y + \frac{1}{2}(\delta^{Nd} - \delta^{Sd})\right)\right), \quad (\text{A9})$$

$$W_{zz}^d(z) = (\delta^{Ud} + \delta^{Dd}) \times \sin\left(\omega_z\left(z + \frac{1}{2}(\delta^{Ud} - \delta^{Dd})\right) + \phi_z\right), \quad (\text{A10})$$

$$W_{tt}^d(t) = (\delta^{Ad} + \delta^{Bd}) \times \cos\left(\omega_t\left(t + \frac{1}{2}(\delta^{Ad} - \delta^{Bd})\right) + \phi_t\right), \quad (\text{A11})$$

and

$$\theta(x, y) = \arctan\left(\frac{y - y_{\text{origin}}}{x - x_{\text{origin}}}\right). \quad (\text{A12})$$

The angle  $\theta(x, y)$  is that defined by the virtual polar coordinate system whose origin is located at  $(x_{\text{origin}}, y_{\text{origin}})$ . The constants  $\omega_z, \omega_t, \phi_z, \phi_t$  are the frequencies and phases

described in relation to  $\mathbf{W}$  and used to fit the weight functions to the data sets. For instance, the value of  $\omega_t$  can be set to  $2\pi/\tau_{\text{total}}$  where  $\tau_{\text{total}}$  is the time length of the heart cycle (in this case the number of cine images for a heart cycle during the acquisition), and  $\phi_t$  would be set to zero if the data acquisition began with early systole (i.e. when the heart was at rest).

The constants  $\lambda^n$  of Equation (A1) are used for keeping the model numerically stable. They are computed using the feature flux control constants  $l^n$  of Subsection 3.4, and the average values  $\bar{w}^d$  of the weights  $W^d$  over their entire range of values:

$$\lambda^n = \frac{l^n}{\sum_{d \in \mathcal{D}} \bar{w}^d} \approx \frac{l^n}{8 \cdot 2/\pi} = \frac{\pi}{16} l^n. \quad (\text{A13})$$