Ultrasound in Med. & Biol., Vol. 27, No. 12, pp. 1583–1594, 2001

0301-5629/01/$–see front matter

PII: S0301-5629(01)00478-1

● *Original Contribution*

ENHANCEMENT OF CONTRAST REGIONS IN SUBOPTIMAL ULTRASOUND IMAGES WITH APPLICATION TO ECHOCARDIOGRAPHY

DJAMAL BOUKERROUI,* J. ALISON NOBLE,* MARC C. ROBINI** and MICHAEL BRADY*
*Medical Vision Laboratory, Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK
**CREATIS, CNRS Research Unit (UMR 5515), INSA de Lyon, FRANCE

(*Received* 6 *June* 2001; *in final form* 2 *October* 2001)

Abstract—In this paper we propose a novel feature-based contrast enhancement approach to enhance the quality of noisy ultrasound (US) images. Our approach uses a phase-based feature detection algorithm, followed by sparse surface interpolation and subsequent nonlinear postprocessing. We first exploited the intensity-invariant property of phase-based acoustic feature detection to select a set of relevant image features in the data. Then, an approximation to the low-frequency components of the sparse set of selected features was obtained using a fast surface interpolation algorithm. Finally, a nonlinear postprocessing step was applied. Results of applying the method to echocardiographic sequences (2-D + T) are presented. The results demonstrate that the method can successfully enhance the intensity of the interesting features in the image. Better balanced contrasted images are obtained, which is important and useful both for manual processing and assessment by a clinician, and for computer analysis of the sequence. An evaluation protocol is proposed in the case of echocardiographic data and quantitative results are presented. We show that the correction is consistent over time and does not introduce any temporal artefacts. (E-mail: djamal@robots.ox.ac.uk)

Key Words: **Image enhancement, Feature detection, Image quality, Phase congruency, Surface interpolation, Echocardiography, Ultrasound.**

INTRODUCTION

The first step toward automatic analysis or evaluation of images frequently consists of feature detection and segmentation. In the field of medical image analysis, robustness, accuracy and reproducibility are critical. Traditionally, in early (low-level) processing and analysis methods, the data are used directly without feature enhancement or bias field (*i.e.*, intensity inhomogeneity) correction. The most important preprocessing step investigated and applied is denoising, which improves the signal-to-noise ratio (SNR). However, denoising is only a partial solution because it cannot correct the image for artefacts introduced by the imaging system. Indeed, it often depends on the imaged object, which finally leads to nonhomogeneous regions in the image. Intensity inhomogeneities are a well-studied problem in the area of image analysis of magnetic resonance imaging (MRI) and ultrasound (US) B-scan images (Hughes and Duck 1997; Xiao et al. 2000). Recently, following the publication of the work by Wells et al. (1996) on bias field correction of MRI data, several authors have investigated this problem and several approaches have been proposed for MRI (Guillemaud and Brady 1997; Velthuizen et al. 1998). Although bias field correction is often necessary for good segmentation, many approaches have exploited the idea that a good segmentation also helps estimation of the bias field (Ahmed et al. 1999; Leemput et al. 1999; Zhang et al. 2001). By contrast, intensity inhomogeneity correction for US images has received relatively little attention, possibly because of the high noise level of B-mode images. Some recent intensity-based adaptive segmentation approaches, which intrinsically take into account the nonuniformity of the tissue classes, have yielded promising results (Ashton and Parker 1995; Boukerroui et al. 1999; Boukerroui 2000). More recently, a novel technique for finding acoustic boundaries

Address correspondence to: Dr. Djamal Boukerroui, Old Library, Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK. djamal@robots.ox.ac.uk

in 2-D and 2-D + T echogram sequences has been proposed (Mulet-Parada 2000; Mulet-Parada and Noble 2000). The most important advantage of this technique is its intensity-independence. However, because the noise rejection in this method involves an intensity-based noise threshold, the method is not truly intensity-invariant and is highly sensitive to noise. The 2-D + T version of the published technique takes advantage of temporal continuity to improve its robustness to noise and detect only relevant and continuous features over time. The authors have reported that spatiotemporal estimation is insufficient for low frame-rate sequences and that there are a number of localisation problems because of the nonuniformity of wall velocity during the cardiac cycle (Mulet-Parada 2000). This underlines the need for development of a feature enhancement approach to correct the image. To our knowledge, the first attempt to adapt bias field correction of Wells et al. (1996) to B-scan US data is the method proposed in (Xiao et al. 2000). Results shown for breast and cardiac US images demonstrate that it can successfully remove intensity inhomogeneities, and significant improvement is achieved in tissue contrast and in the resulting image segmentation. The approach is promising. However, as reported in (Xiao et al. 2000), it still requires user interaction to set the image model parameters.

In this paper, we propose a novel feature-based contrast enhancement approach. Our approach is based on a phase-based feature detection algorithm, followed by sparse surface interpolation and subsequent nonlinear postprocessing. We first exploited the intensity-invariant property of phase-based acoustic feature detection to select the relevant features in the data. Then, an approximation to the low-frequency components of the sparse set of selected features was obtained, using a surface interpolation algorithm. Finally, a nonlinear postprocessing step with one control parameter was applied.

The paper is organised as follows. In the Methods section, we briefly describe the mathematical framework underlying the sparse surface interpolation algorithm (Yaou and Chang 1994), review the feature asymmetry (FA) measure for 2-D acoustic boundary detection (Mulet-Parada 2000), and then present the nonlinear postprocessing stage. The proposed evaluation protocol in the special case of echocardiographic data, as well as quantitative results using the approach, are presented and discussed in the Results section. The authors' conclusions are summarised in the Conclusion section.

METHODS

Overview of algorithm

A block diagram of our new feature-based contrast enhancement method is shown in Fig. 1. First, features are detected in the image $I(x, y)$. This provides a normalised likelihood image $FA_{2\text{-}D}(x, y)$ where the intensity value at a position (x, y) is proportional to the significance of the detected features. The $FA_{2\text{-}D}(x, y)$ measure varies from a maximum of 1 (indicating a very significant feature) down to 0 (indicating no significance). The feature detector that we used is based on phase congruency (PC) (Morrone and Owens 1987; Kovesi 1999) because it provides a single unified theory that is capable of detecting a wide range of features, rather than being specialised for a single feature type, such as an intensity step edge. Inspection of cardiac B-mode US images reveals features of a number of quite distinctive kinds; for example, the intensity ridge corresponding to (a large part of) the epicardial wall and the mitral valve. Further, PC is theoretically invariant to brightness and contrast. Hence it is, in principle, robust against typical variations in image formation. The disadvantage of PC is a direct result of contrast invariance; namely, its high sensitivity to noise. The poor SNR of B-mode US, including heavy speckle, means that this is a problem that has particularly to be addressed when applying PC to cardiac US.

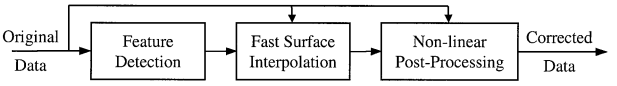

Fig. 1. Block diagram of the feature enhancement method.

Following feature detection, the sparse data at feature locations are then interpolated by a fast sparse surface interpolation algorithm using the likelihoods to estimate the degradation field; then, a novel nonlinear processing method using the degradation field is applied to the original data to enhance or de-emphasise feature values. These steps are described next.

2-D sparse surface interpolation

In this section, we review the method we employ for fast sparse surface interpolation. This is based on Yaou and Chang (1994), to which the reader is referred for further details.

Surface interpolation from a sparse set of noisy measured points is an ill-posed problem because an infinite set of surfaces can satisfy any given set of constraints. Hence, a regularisation procedure, taking into account the visual relevance and computational efficiency is usually applied, so that the interpolation problem becomes a minimisation of an energy functional of the form:

$$U(f) = U_d(f, d) + \lambda U_r(f). \quad (1)$$

The first term in eqn (1) is a faithfulness measure of the approximating surface f to the measured data d and is called the "cost or constraint functional." The second is the regularisation functional; λ is a nonnegative (Lagrange multi-

plier) parameter controlling the amount to which the data are to be considered (piecewise) smooth. A commonly used cost functional is the weighted sum of squares:

$$U_d(f, d) = \sum_{i=1}^{n} w_i(f(x_i, y_i) - d_i)^2, \tag{2}$$

which measures the difference between the measured field $d = \{(x_i, y_i, d_i)\}$ and the approximating surface $f(x_i,y_i)$; and where $w = \{0 \le w_i \le 1\}$ is the corresponding set of weights for the measured field, reflecting the confidence of the measured information at each position ($w_i = 0$ means the absence of information at (x_i, y_i)); n being the number of data points. Regarding the regularisation term, a common approach has been to use a variational functional to constrain the solution, often expressed as a thin plate energy, otherwise known as the quadratic variation (Grimson,1981):

$$U_r(f) = \int\int \left[\left(\frac{\partial^2 f}{\partial x^2}\right)^2 + 2\left(\frac{\partial^2 f}{\partial x \partial y}\right)^2 + \left(\frac{\partial^2 f}{\partial y^2}\right)^2\right] \mathrm{d}x\mathrm{d}y. \tag{3}$$

In general, obtaining an analytic solution of the Euler–Lagrange equations resulting from the above optimisation problem is difficult. Therefore, an approximation to the continuous problem using discrete operators is used. This leads to a numerical solution. Suppose that the data d is defined on a regular rectangular lattice $G = \{(x_i, y_j), 1 \le i, j \le N\}$, and that a discrete representation of the surface is defined using a set of nodal variables $\nu = \{\nu_{i,j} = f(x_i, y_j)\}$. The discrete representation of the cost functional in eqn (2) is:

$$U_d(\nu, d) = \sum_{i,j} w_{i,j}(\nu_{i,j} - d_{i,j})^2. \tag{4}$$

By concatenating all the nodal variables $\nu_{i,j}$ and the data $d_{i,j}$, respectively, into column vectors $\boldsymbol{\nu}$ and $\boldsymbol{d}$, we obtain the usual matrix representation of eqn (4),

$$U_d(\nu, d) = (\boldsymbol{\nu} - \mathbf{d})^{\mathrm{T}} \boldsymbol{A}_w (\boldsymbol{\nu} - \boldsymbol{d}), \qquad \boldsymbol{A}_w = \mathrm{diag}(\{w_{i,j}\}). \tag{5}$$

Regarding the regularisation term, the finite element method provides a continuous surface approximation, which is a good means of converting the continuous expression for the energy and leads to a tractable discrete problem that has a numerical solution. The discrete form of the thin plate is given by:

$$U_r(\nu) = \sum_{i,j} \left\{ \begin{array}{c} (\nu_{i+1,j} - 2\nu_{i,j} + \nu_{i-1,j})^2 \\ +2(\nu_{i+1,j+1} - \nu_{i,j+1} - \nu_{i+1,j} + \nu_{i,j})^2 \\ + (\nu_{i,j+1} - 2\nu_{i,j} + \nu_{i,j-1})^2 \end{array} \right\}, \tag{6}$$

which can be reorganised as follows:

$$U_r(\nu) = \sum_{i,j} \sum_{m,n} \nu_{i,j} a_{i,j,m,n} \nu_{m,n}, \tag{7}$$

where the coefficients $a_{i,j,m,n}$ describe the relations between the nodal variables $\nu_{i,j}$ and $\nu_{m,n}$ and are called relation coefficients. Hence, the discrete version of the continuous regularisation term given by eqn (3) is:

$$U_r(\nu) \propto \boldsymbol{\nu}^{\mathrm{T}} (\lambda \boldsymbol{A}_r) \boldsymbol{\nu}, \tag{8}$$

where $\boldsymbol{A}_r$ is $N^2 \times N^2$ sparse matrix and contains at most 13 nonzero elements per line. Finally, by adding eqns (5) and (8), we obtain the corresponding discrete version of the functional in eqn (1):

$$U(\nu) = \boldsymbol{\nu}^{\mathrm{T}} \boldsymbol{A} \boldsymbol{\nu} - 2\boldsymbol{\nu}^{\mathrm{T}} \boldsymbol{b} + c, \tag{9}$$

where, $\boldsymbol{A} = \lambda \boldsymbol{A}_r + \boldsymbol{A}_w$ and $\boldsymbol{b} = \boldsymbol{A}_w \boldsymbol{d}$, and c is a constant. The resulting energy function has a minimum at $\boldsymbol{\nu} = \boldsymbol{\nu}^*$, the solution of the linear system $\boldsymbol{A}\boldsymbol{\nu} = \boldsymbol{b}$, with a very sparse system matrix $\boldsymbol{A}$.

Resolution in the wavelets space

The surface interpolation problem described in the last section leads to the solution of a large linear system with a sparse system matrix $\boldsymbol{A}$. Therefore, the equation system is nearly singular and results in poor convergence when simple iterative methods are used. To obtain fast surface interpolation, a scheme is needed that can improve the numerical conditioning. The multigrid (Terzopoulos 1983) and hierarchical basis techniques (Szeliski 1990) have been applied successfully to speed up convergence. Both of these techniques use a multiresolution approach. However, in both cases, the spatial frequency domain property of the interpolation problem is not effectively utilized. Recently, a more tractable approach, in terms of simplicity and efficiency, has been proposed (Yaou and Chang 1994). The approach utilizes the concept of preconditioning in a wavelet transform space. Preconditioning techniques aim to reduce the condition number κ of the system matrix (defined as the ratio between the magnitude of the largest and the smallest eigenvalues), which implies speeding up the rate of convergence. Wavelet transform analysis provides a powerful way to simultaneously filter data in space and scale and is easily implemented using filter banks (Mallat

1999). Yaou and Chang (1994) use vector space decomposition reconstruction of the wavelet transform to analyse the multiresolution representation of the surface interpolation problem. The authors exploit the fact that, usually, the high-frequency component of the interpolation problem converges much faster than the low-frequency component in surface interpolation. Interestingly, this is the same observation that inspired Terzopoulos' use of the multigrid algorithm. The use of the discrete wavelet transform (DWT) allows the low-frequency component and the high-frequency component of the interpolation problem to be solved separately.

In other words, the minimisation is carried out in a wavelet space using an asynchronous iterative computation and a biorthogonal spline wavelet basis for the preconditioning step (Yaou and Chang 1994). The DWT preconditioning transfers the equation system to an equivalent one with new nodal variables $\tilde{\boldsymbol{v}}$ and a new system matrix, $\tilde{\boldsymbol{A}}$, that is much denser than the original one, $\boldsymbol{A}$. This implies that a more global connection between the interpolation nodes can be made which considerably improves the convergence rate of the iterative solution. To give the reader some insight on the relative convergence rate with and without preconditioning in our problem, for a 64×64 interpolation problem, the value of κ is reduced by a factor of 2.5 and 5.4, respectively, when two and three decomposition levels for the DWT are used. We refer the reader to Yaou and Chang (1994) for more details about fast surface interpolation using the multiresolution wavelet transform, and to Cohen et al. (1992), Daubechies (1992), and Mallat (1999) regarding the biorthogonal wavelet basis and the associated filter banks.

Phase-based feature detection

Phase-based feature detection has been investigated extensively following the publication of the local energy model of feature detection (Morrone and Owens 1987). This model postulates that features can be defined and classified using their phase signatures, or their phase congruency (PC). In other words, feature information is encoded at points where all the Fourier components are maximally in phase. Thus, the PC is maximum. These observations have led to the development of a number of phase-based feature detection algorithms (Kovesi 1999; Mulet-Parada 2000; Mulet-Parada and Noble 2000) (and references therein). In particular, measures based on phase information seem to be more appropriate for acoustic feature detection because US images are characterized by a low SNR due to the presence of speckle and to the high range of imaging artefacts, causing the alteration of the intensity magnitude of equally significant features in the data.

Strictly speaking, the concept of PC is only defined in one dimension because its definition involves the Hilbert transform. Typically, the computation of PC (and the related concept of local energy) uses a pair of quadrature filters (Venkatesh and Owens 1990), normally log-Gabor filters. A series of oriented 2-D filters can, for example, be constructed by 'spreading' a log-Gabor function into 2-D. In this way, extension to 2-D of the 1-D phase measure may be obtained (Kovesi 1999).

In our work, we have used the 2-D feature asymmetry (FA) measure used in Mulet-Parada (2000), and Mulet-Parada and Noble (2000) for feature detection. This measure provides good detection of asymmetric image features, such as step edges, and has the advantage of being intensity invariant. The 2-D FA measure is defined by:

$$FA_{2\text{-}D}(x, y) = \sum_m \frac{\lfloor |o_m(x, y)| - |e_m(x, y)| - T_m \rfloor}{\sqrt{o_m^2(x, y) + e_m^2(x, y)} + \varepsilon}, \tag{10}$$

which is a sum over m orientations of a normalised measure of the difference between the odd $o_m(x,y)$ and the even $e_m(x, y)$ filter responses. Here, $\lfloor\ \rfloor$ denotes zeroing of negative values, ε is a small positive number to avoid division by zero and T_m is an orientation-dependent noise threshold, defined by (Mulet-Parada 2000):

$$T_m = k \cdot std[|o_m(x, y)| - |e_m(x, y)|], \tag{11}$$

where k is a positive factor controlling the noise threshold and *std* is the standard deviation (SD). For more details about the implementation and the setting parameters of the log-Gabor filters and the spreading functions, see Kovesi (1999), Mulet-Parada (2000) and Mulet-Parada and Noble (2000).

The output of applying the FA algorithm is a location image, which is a normalised likelihood image [0, 1], where the intensity values can be interpreted as a confidence measure of feature detection.

The new feature-enhancement algorithm

Having reviewed the two principal algorithms that are used in our new feature detection method, we now address how these are used for feature enhancement. Briefly, our method involves reconstructing an approximation to the intensity inhomogeneities that can be utilised to correct the original corrupted data. A mathematical model for the intensity inhomogeneity in US images was developed by Hughes and Duck (1997). The authors used a multiplicative degradation model. Motivated by this, we define a correction equation as:

$$I_c(x, y) = \frac{I(x, y)/\max(I(x, y))}{\nu^*(x, y)/\max(\nu^*(x, y)) + \gamma}. \quad (12)$$

Here, $\nu^*(x, y)$ is the interpolated surface and γ is a positive control parameter that ensures that $I_c(x, y) \propto I(x, y)$ for $\gamma >> 1$. The maximum correction is obtained when $\gamma << 1$. Assuming that the image intensities of occurrences of a single tissue type should be equal, an estimate of the low frequency components of an intensity data field can be made by taking the image intensity values of the image only at the locations of the relevant features. An estimate of the base frequency of this degradation can be found using the fast surface interpolation algorithm as follows. We define the set of nodal variables ν and the corresponding weighting field w, by:

$$\nu = \{\nu_{i,j} = \max_{B_{i,j}} I(x, y) \text{ if } FA_{2\text{-}D}(x_i, y_j) > 0; \quad 1 \le i, j \le N\} \quad (13)$$

$$w = \{w_{i,j} = FA_{2\text{-}D}(x_i, y_j);\ 1 \le i, j \le N\}. \quad (14)$$

Here, $B_{i,j}$ is a small rectangular window centred at pixel position (x_i, y_j) and N is the size of the data. Taking the maximum intensity value in a window centred on the feature position guarantees that we always take the highest value of the step edge. Hence, eqns 13 and 14 define the two inputs for the surface interpolation stage, which provide us with an estimate of the low-frequency image degradation.

Implementation details

In this section, we present some implementation details related to our method and its efficiency. The surface interpolation problem is difficult because the nodal variables are defined on a very sparse grid. For instance, in the particular case of echocardiography images, the nonexistence of any constraint variables (nodes) in the left ventricular cavity causes problems at the end diastolic period of the cardiac cycle when the area of the cavity is very large. On the other hand, derivation of the system matrix $\tilde{\boldsymbol{A}}$ requires $O(n^2)$ operations where n is the number of the nodal variables, so we aim to keep n small. To avoid these two problems, an estimation of the surface using low-pass filtered and subsampled data is carried out; then we interpolate the resulting surface on a regular grid to provide the estimate of the bias field of the original data. We found that the use of low-pass downsampled data benefits both the fast surface interpolation stage and the 2-D FA detection stage, because the filtering operation allows reduction of the speckle noise. Finally, a histogram threshold is found to eliminate the high values due to the enhancement of noise peaks. The threshold h, is defined empirically by the following equation:

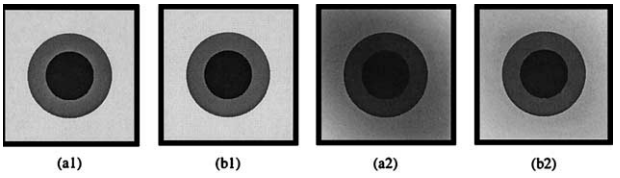

Fig. 2. Original images: (a1) Ideal; (a2) corrupted, and their corresponding enhanced images (b1) and (b2). Results obtained for $N = 64$ and $\gamma = 0.2$.

$$P_{I_c}(I_c(x, y) < h) = 0.999. \quad (15)$$

Here, $P_{Ic}(\cdot)$ is a probability measure. The corrected image is rescaled between [0 to 255] and the grey value of the rejected points is set to 255 (see Fig. 2c1). We found the rejection of a small number of points sufficient to cut the tail of the distribution without introducing false edges or modifying the shape of the distribution. Cutting the tail of the distribution (which contains no information) has the advantage of enlarging the intensity range of the corrected data.

Note that the nonlinear processing given by eqn (12) and the thresholding operation are applied to the whole sequence when the processed data are 2-D + T.

RESULTS AND QUANTITATIVE EVALUATION

All the results presented in this paper are obtained for $k = 0.45$ (FA threshold factor), $\lambda = 1$ (FSI parameter) and with 2 decomposition levels for the DWT in the fast surface interpolation algorithm.

Synthetic data

The aim of the first experiment was to show, for synthetic data, that the proposed approach is capable of removing (or at least reducing) the bias field without introducing any artefacts. Figure 2 shows two images, an ideal one (a1) and corrupted one (a2) and their corrections (b1 and b2, respectively). We can see that a significant contrast enhancement is obtained in both cases and that the corrected images appear visually similar.

Real data

One short-axis and two long-axis sequences of approximately 100 frames were chosen randomly from the ADEQUATE project database of routine acquisition scans at the Hospital General Universitario Gregorio Marañón, Madrid and at the John Radcliffe Hospital, Oxford. The data sets were acquired on an HP-SONOS

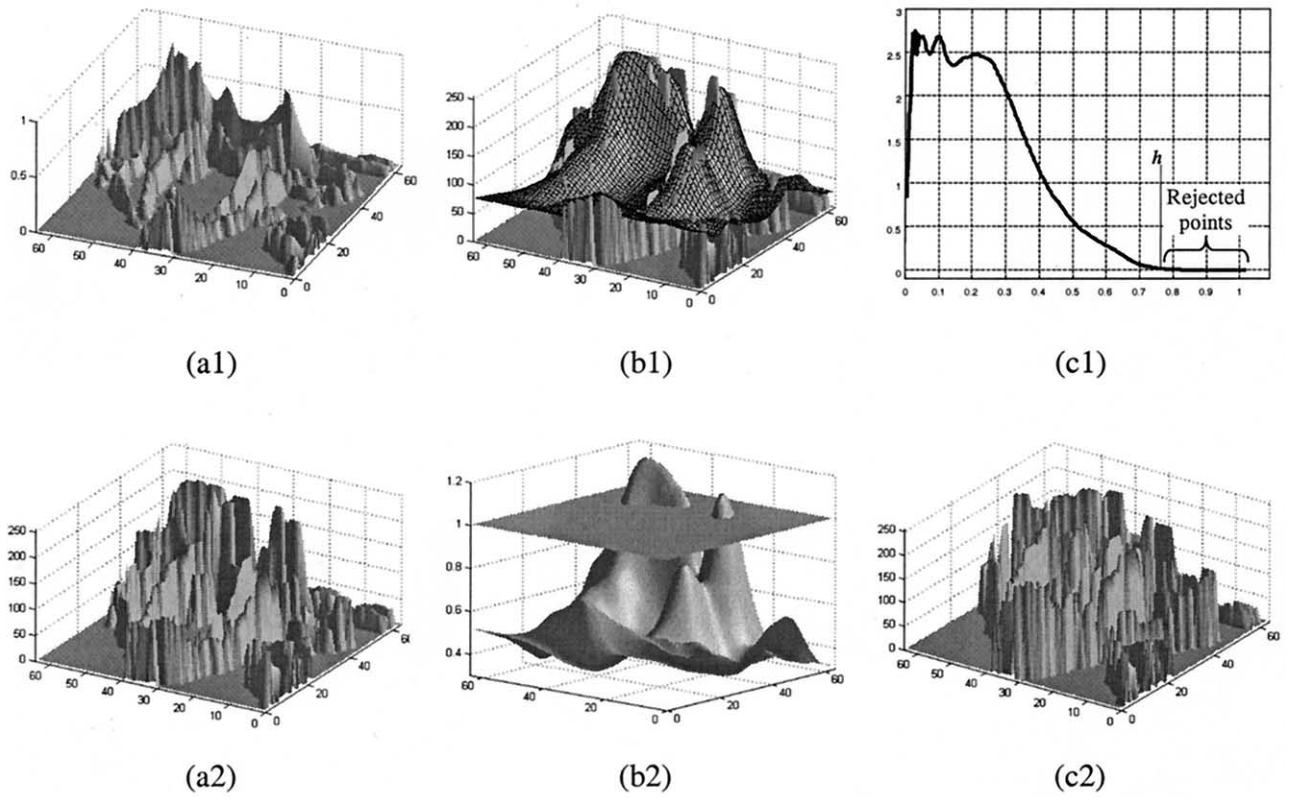

Fig. 3. Illustration of steps in the algorithm. (a1) Likelihood image representing the weighting field w, for the surface interpolation stage; (a2) the original data at the location of the detected boundaries representing the data field ν, for the surface interpolation stage; (b1) the result of surface interpolation superimposed on the data field; (b2) the normalised interpolated surface with the additional shift, $\gamma = 0.2$, showing clearly the region where the intensity will be lowered and the region where it will be enhanced; (c1) histogram thresholding operation; (c2) corresponds to (a2), but shows it after correction.

5500 scanner (Agilent Technologies). Data sets 1 and 3 were fundamental mode acquired at a frame rate of approximately 80 Hz and transducer frequency of 3.6 MHz. Data set 2 was acquired at a frame rate of 25 Hz, using second harmonic image enhancement (transmit/receive 1.8/3.6 MHz). The size of the original images was 720 $\times$ 512 pixels. A region-of-interest (ROI) of 512 $\times$ 512 in the centre of each frame was extracted.

Figure 3 shows the original data and results at intermediate stages of processing for an echocardiographic image. The image used in this experiment is shown in Fig. 4a1. Notice the correlation between the likelihood image (a1) and the intensity image of the detected features (a2). Hence, if the SNR is low, the 2-D FA measure does not yield a clean feature detection image. In that case, either the noise threshold, eqn (11) has to be set to a higher value, which increases the false-negative detection rate, or it has to be set to a low value, in which case the false-positive rate will increase. Comparison of the images (a2) and (c2) provides some (qualitative) insight about how much the features have been enhanced.

To understand the effect of using approximate data at a coarse resolution, we compared feature enhancements using different sampling resolutions. As illustrated by the example shown in Fig. 4, we have found that we do not observe any significant difference between the results obtained using approximations to the original data (512 $\times$ 512) on a 64 $\times$ 64 or a 128 $\times$ 128 grid. Both corrections seem to give acceptable corrections and contrast enhancement (see the profile lines). However, the plot of the vertical line profile Fig. 4c3 shows that the background signal (interior of the LV cavity) is more enhanced for $N = 128$ than for $N = 64$. This is due to the problem of sparse surface interpolation mentioned in the Implementation Details section. Figure 4b1, b2, b3 show the results of the 2-D FA boundary detection on the original image and on the enhanced images. Significant improvements in boundary detection on the enhanced images is observed, particularly in the apex region where the intensity values of the original image are very low.

In the third experiment, we studied the influence of the control parameter γ, eqn (12). Figure 5 shows the

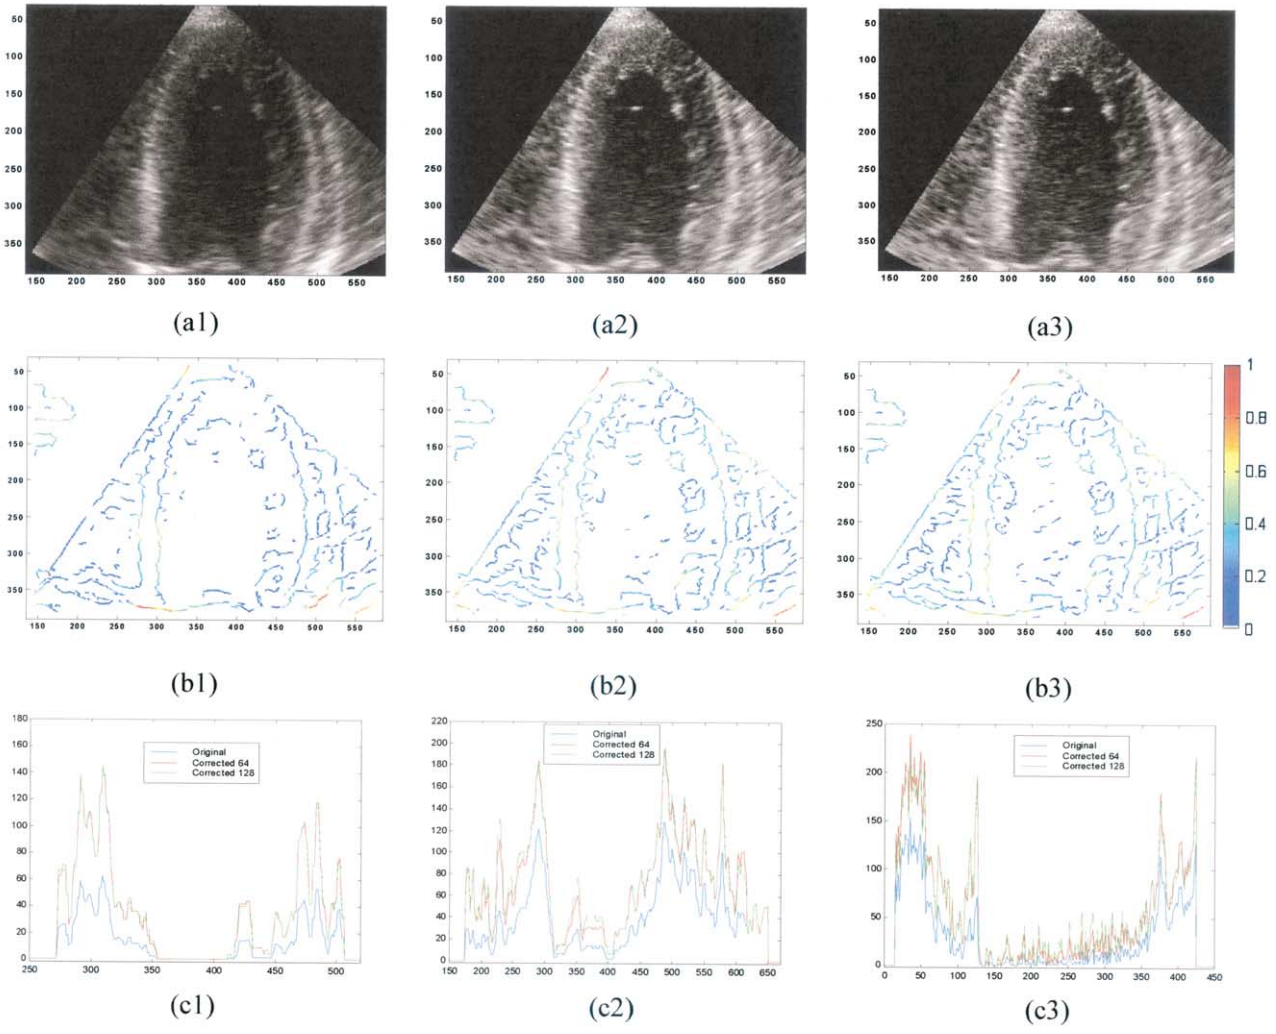

Fig. 4. Frame 11 of data set 1. Comparison of (a1) the original image and (a2) the enhanced one obtained with $N = 64$ and (a3) $N = 128$. Images (b) show the corresponding feature asymmetry results (likelihood image). Figures (c) show profile lines: (c1) horizontal line 150; (c2) horizontal line 350; and (c3) vertical line 355. Results obtained with $\gamma = 0.2$.

enhancement results and the corresponding feature asymmetry detection on the corrected image for three different values of γ: 0.4, 0.2 and 0.05. The three results are good and appear, at first sight, to be similar. However, on more careful observation, we note that the enhancement increases when γ decreases. The plots of line profiles shown in Fig. 6 offer a clear demonstration of the influence of γ. Notice that the three results appear to be in good agreement where the signal is high and low signal values are more enhanced for $\gamma = 0.05$ than 0.2 or 0.4. However, this observation does not mean that the enhancement result for $\gamma = 0.05$ (or $\gamma < 0.05$) is better than for the other two. Indeed, if $\gamma = 0.05$ enhances the low signal values better than $\gamma = 0.2$ (0.4), it does the opposite for high signal values. An objective evaluation and quantitative results of the enhancement are necessary to answer the question as to which values of γ give the best enhancement. This is described next. Figure 7 shows the enhancement results and the corresponding feature asymmetry detection for the original and the corrected images for the two other data sets (second harmonic long axis and B-mode short axis).

Quantitative evaluation. A standard protocol in the evaluation of heart disease is to assess regional and global left ventricular function, using quantitative and qualitative measurement based on the myocardial walls (epicardial and endocardial). This includes heart wall delineation, wall thickness estimation and motion estimation, like muscle contraction and the opening and closing of the cardiac valves. Clinicians desire better image quality. The automatic analysis of echocardiograms is, to date, a challenge due to the low contrast and very noisy appearance of standard grey-level (B-mode) images. The quantitative evalua-

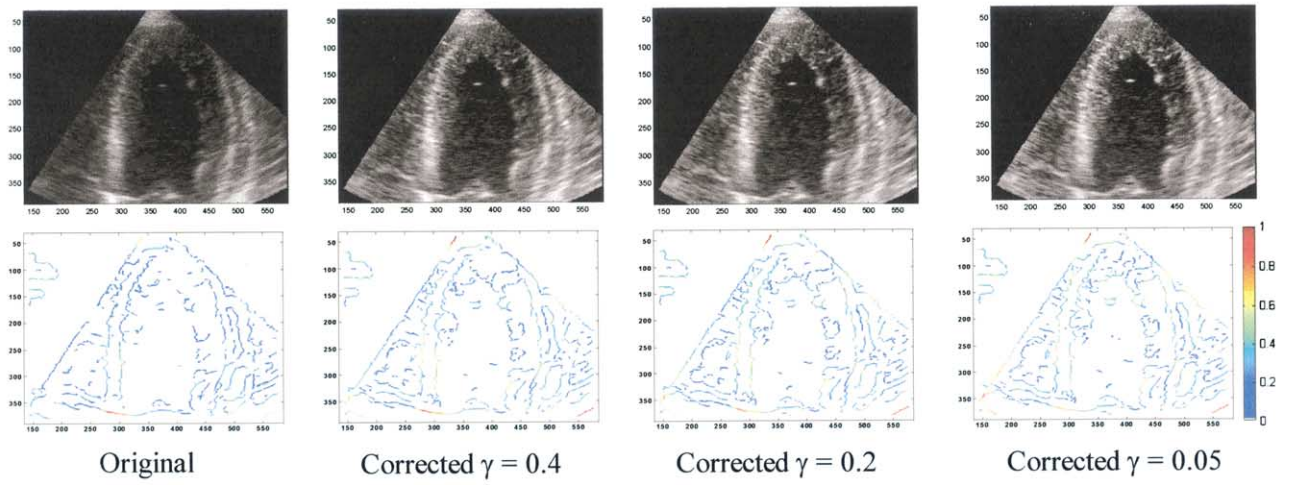

Fig. 5. Influence of the control parameter γ. The original images correspond to frame 13 of data set 1. Results obtained for $N = 64$.

tion of the proposed feature-based contrast enhancement approach will be described next.

Quantitative evaluation of computer vision algorithms is an important task, particularly in the case of medical imaging. The availability of ground truth makes this task easier. However, there is no ground truth for the data available for the current study. To overcome this problem, we have to answer the question: Which are the important features in the data and how can an automatic detection algorithm respond to the features and not others? In the case of echocardiographic images, and for the purpose of this paper, we are interested in the detection of the endocardial boundary. Because these features are often modelled as step edges, a measure of the height of the step is a good evaluation parameter. A second criterion that our enhancement should satisfy is consistency over time (*i.e.*, the method should not introduce temporal artefacts). A good measure of temporal consistency is the correlation coefficient between measures calculated on the enhanced images and the same measures calculated on the original data. We investigate both aspects next.

For each image, we define three regions next to the interesting features: RC is the region located in the cardiac cavity near to the endocardial wall, RM is the myocardium, and RE is the region located from the epicardial border outwards. Briefly, for each image, seven segments were defined by selecting points on a manually defined interpolating spline fit to the endocardial border. The intensity profiles crossing the three regions (RE, RM, RC) could then be automatically defined by computing the normal to the curve and sampling pixels along the normal direction, as defined before. The width of each region was about five pixels. This is illustrated in Fig. 8. For each of the seven segments, the

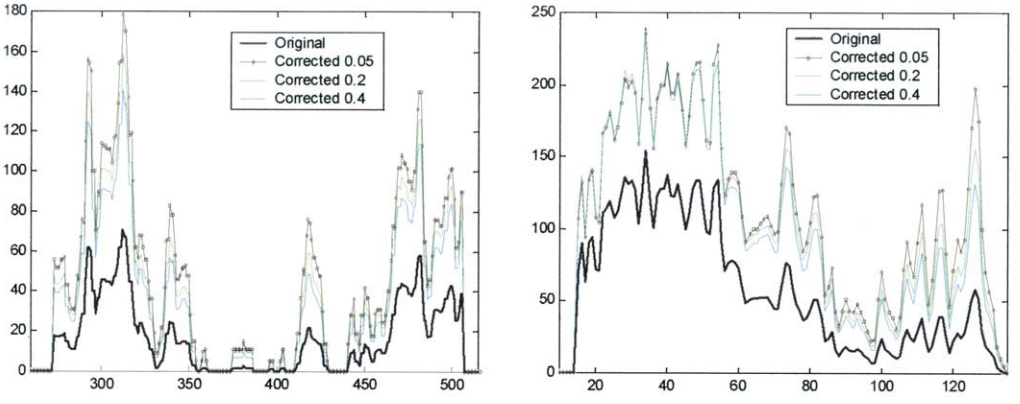

Fig. 6. Influence of γ. Line profiles of the results shown in Fig. 5. Left: Horizontal line 150; right: vertical line 355. Observe, on the right side of the figure, the enhancement of the peak corresponding to the endocardium border on the apex.

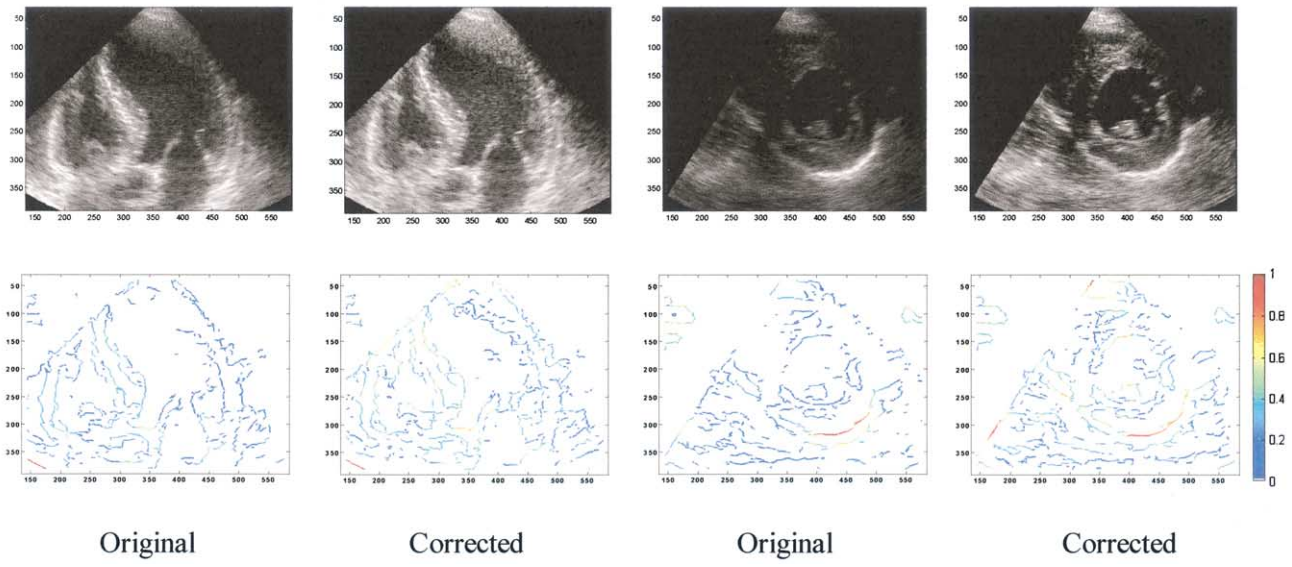

Fig. 7. Correction results on two other data sets. Left: data set 2, frame 17; right: data set 3, frame 0. Results obtained for $N = 64$ and $\gamma = 0.2$.

contour was sampled at 600 points to generate 600 intensity profiles. In practice, we tried to avoid the papillary muscles in the cavity. In the case of the short-axis images, it was difficult to know where the epicardial border was, so we did not measure RE.

For each image, we computed ten measures: mean (signal) and SD of the cavity signal (RC), of myocardium signal (RM), of the signal beyond the epicardial border (RE) and of the differences (RE − RM and RM − RC). We then computed the mean and the SD (μ_t, σ_t) of each of these measures over time. From this, we could investigate if there was a significant contrast improvement, especially at features location, namely the cardiac walls (see RE − RM and RM − RC values).

Table 1 shows the computed measures on a long axis data set for the original images and those enhanced for $N = 64$ and $N = 128$. Note the similarity of the calculated measures and the small relative differences, demonstrating the nonexistence of significant difference between the two results. This table confirms the conclusion of experiment 2 because the highest relative difference is for the calculated measures for the cardiac cavity region (RC) (the cavity signal is more enhanced for $N = 128$).

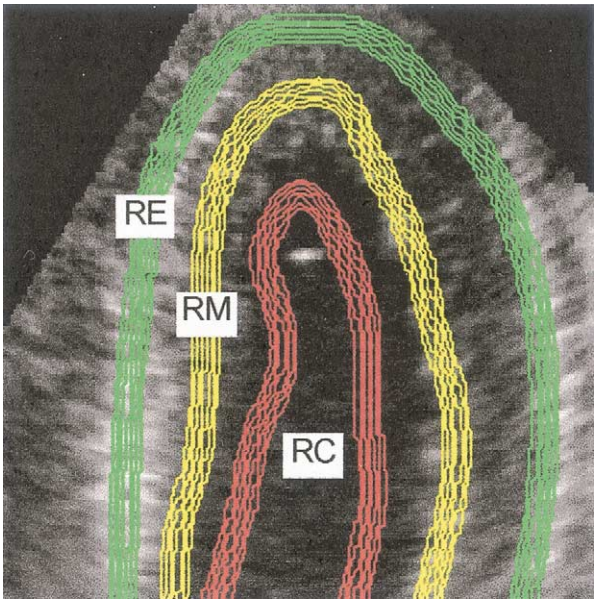

Fig. 8. Illustration of region-of-interest for the evaluation protocol.

Table 1. Quantitative comparison of the enhancement for two different sampling resolutions (64 and 128)

Data set 1 Long axis		Original (μ_t, σ_t)	Corrected 64 (μ_t, σ_t)	Corrected 128 (μ_t, σ_t)
RC	Signal	6.81, 1.02	16.37, 2.16	18.88, 2.57
	SD	6.94, 1.14	13.07, 2.22	15.59, 2.65
RM	Signal	29.70, 3.29	61.86, 5.45	64.01, 5.47
	SD	18.61, 2.21	29.62, 2.42	30.38, 2.38
RE	Signal	61.18, 3.99	107.64, 5.38	107.07, 5.33
	SD	35.53, 3.20	44.71, 3.59	45.14, 3.87
RE−RM	Signal	31.48, 4.12	45.78, 6.10	43.06, 6.33
	SD	28.59, 2.63	44.96, 3.12	44.88, 3.11
RM−RC	Signal	22.89, 3.20	45.49, 5.70	45.14, 5.64
	SD	15.98, 1.77	28.34, 2.35	29.13, 2.37

Here, signal is the spatial mean of the signal. The table shows the means and the SDs over the frames of one cardiac cycle. $\gamma = 0.2$.

Table 2. Quantitative analysis of the enhancement and study of the influence of the control parameter γ on the enhancement for the three data sets

		Original (μ_t, σ_t)	Corrected 0.4 (μ_t, σ_t)	Corrected 0.2 (μ_t, σ_t)	Corrected 0.1 (μ_t, σ_t)	Corrected 0.05 (μ_t, σ_t)
Data set 1 Long axis						
RC	Signal	6.81, 1.02	14.80, 1.94	16.37, 2.16	17.65, 2.32	18.49, 2.42
	SD	6.94, 1.14	12.17, 2.02	13.07, 2.22	13.59, 2.36	13.90, 2.44
RM	Signal	29.70, 3.29	56.16, 5.23	61.86, 5.45	66.09, 5.51	68.84, 5.48
	SD	18.61, 2.21	28.51, 2.64	29.62, 2.42	30.29, 2.19	30.86, 2.07
RE	Signal	61.18, 3.99	101.71, 5.33	107.64, 5.38	110.89, 5.43	112.40, 5.47
	SD	35.53, 3.20	45.93, 3.65	44.71, 3.59	42.98, 3.53	41.59, 3.49
RE−RM	Signal	31.48, 4.12	45.54, 5.84	45.78, 6.10	44.81, 6.31	43.56, 6.48
	SD	28.59, 2.63	42.67, 3.15	44.96, 3.12	46.72, 3.13	48.03, 3.23
RM−RC	Signal	22.89, 3.20	41.36, 5.33	45.49, 5.70	48.43, 5.85	50.35, 5.90
	SD	15.98, 1.77	26.20, 2.26	28.34, 2.35	30.18, 2.52	31.66, 2.74
Data set 2 Long axis						
RC	Signal	6.45, 0.64	13.88, 1.12	15.25, 1.20	16.39, 1.27	17.11, 1.33
	SD	6.63, 0.67	11.20, 0.97	11.89, 1.05	12.26, 1.12	12.46, 1.19
RM	Signal	30.57, 3.54	57.52, 5.64	63.17, 5.92	67.24, 6.05	69.81, 6.11
	SD	19.16, 2.18	29.36, 2.87	30.39, 2.80	30.84, 2.68	31.09, 2.59
RE	Signal	61.87, 3.78	102.56, 4.89	108.39, 4.91	111.52, 4.99	112.90, 5.10
	SD	35.67, 2.49	46.18, 3.11	45.03, 3.23	43.35, 3.32	42.01, 3.40
RE−RM	Signal	31.29, 4.63	45.04, 6.50	45.22, 6.75	44.27, 6.97	43.09, 7.14
	SD	28.19, 2.35	42.03, 2.64	44.24, 2.70	45.86, 2.82	46.99, 3.01
RM−RC	Signal	24.12, 3.49	43.63, 5.63	47.92, 5.98	50.86, 6.18	52.69, 6.29
	SD	16.42, 2.03	26.48, 2.84	28.29, 2.85	29.66, 2.84	30.64, 2.86
Data set 3 Short axis						
RC	Signal	1.41, 0.62	4.31, 1.64	5.05, 1.95	5.68, 2.14	5.53, 2.09
	SD	2.05, 0.79	5.59, 1.65	6.56, 1.94	7.36, 2.14	7.18, 2.10
RM	Signal	18.55, 3.40	43.17, 8.90	50.76, 11.12	56.28, 12.95	54.78, 12.99
	SD	12.04, 1.21	24.94, 2.11	28.85, 2.61	31.72, 3.25	30.72, 3.47
RM−RC	Signal	17.14, 3.01	38.86, 7.74	45.71, 9.72	50.60, 11.39	49.25, 11.46
	SD	12.11, 1.33	25.14, 2.13	29.02, 2.30	31.86, 2.57	30.80, 2.63

Here, signal is the spatial mean of the signal. The table shows the means and the SDs over the frames. $N = 64$.

Table 2 presents the evaluation measures for the third experiment (influence of γ). Note the small values of the SD over time for all the computed measures, both for the original and the enhanced images, indicating a good time consistency. We observe that the signal enhancement is proportional to γ^{-1} and the SDs for RC and RM increase slightly. However, the RE SD decreases with γ^{-1}. This is because, as we have already noted, the highest grey-level values will be reduced, and the lower grey-level intensities will be increased (see Fig. 3b2). When the enhancement is high (γ small), a "saturation phenomenon" appears at the highest intensity values. Because the RE region corresponds to the highest grey values in the images, the spatial SD of this region will decrease with enhancement.

Analysis of the signal differences reveals that both the signal and the SD of the difference RM − RC increase as γ decreases. However, the signal of the difference RE − RM increases as γ decreases until $\gamma = 0.2$ and decreases for $\gamma > 0.1$. This is a consequence of the "saturation phenomenon" as the step edge RE − RM is at high intensities.

These observations enable us to understand more fully the behaviour of the enhancement in function of the parameter γ. The reader may notice how close the quantitative measures are for the different values of the control parameter. In our experiments, we found that a value between 0.1 and 0.2 gives acceptable enhancement results and good trade-offs between the enhancement of the low (RM − RC) and high (RE − RM) grey-level intensity features.

Consistency. Because two key parts of our feature enhancement algorithm do not take into account the temporal information, the consistency of the enhancement over time needs to be studied to show that no temporal artefacts are introduced. Figure 9 shows the spatial correlation curves (where correlation has been computed between consecutive frames) for the original sequences and for their corresponding enhanced sequences. In other words, each curve shows the temporal evolution of the spatial correlation coefficient of consecutive frames in the sequence. The interesting aspect of these curves is not the absolute values of the spatial correlation, but its evolution over the frames. If the enhancement method is consistent over time, the evolution of the curves over time should be the same. Indeed, the curves show that the temporal correlation of the

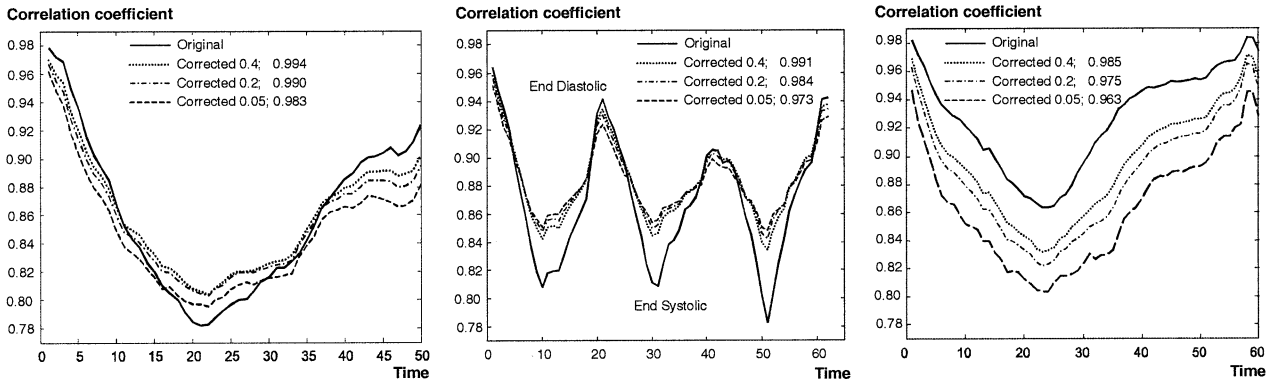

Fig. 9. Evolution over time of the spatial correlation coefficients between consecutive frames. Left: Data set 1; middle: data set 2; right: data set 3. "Corrected 0.2; 0.990" means results obtained for $\gamma = 0.2$ and the correlation coefficient of the curve to the equivalent curve of the original data is 0.990. $N = 64$.

original data is well conserved in the corrected sequences (see the temporal correlation coefficients in the figure legend). Notice that the temporal continuity is better conserved for the two first data sets (long axis) than for the third one (short axis). The lowest correlation coefficient for the temporal continuity is 0.963 obtained for the enhancement of data set 3 with $\gamma = 0.05$.

The second experiment analysing temporal consistency involved computing the spatial correlation coefficient between the enhanced and the original data frame-by-frame. If the enhancement method is consistent over time, the value of the spatial correlation should be almost constant for all the frames. In other words, we expect the correlation coefficient to be independent of time. Table 3 shows the mean and the SD over the time of this measure for the three processed data sets. Notice how small the SDs are, indicating good temporal consistency of the enhancement.

A comparison between the results of the three data sets gives us some insight into the robustness of the proposed approach.

Finally, Fig. 10 shows further enhancement results obtained on four different data sets (one frame each). The grey-level data were acquired in the same condition as "before" data. All examples are typical of those in clinical practice (*i.e.*, were not specifically acquired for this paper, but were routine scans). The last example (right image) was acquired using a contrast enhanced power Doppler (CEPD) imaging system. CEPD detects the contrast agent in the cardiac cavity that has been previously injected into the patient. The enhancement results suggest that the proposed method could be used as a preprocessing for CEPD images because it reduces the intensity inhomogeneity of the signal in the cavity region. This potentially suggests that our method may be very useful for improving automatic detection of the endocardial border. This is currently being investigated in the ADEQUATE project.

Table 3. Evaluation experiment of the consistency over time

	$\gamma = 0.4$; (μ_t, σ_t)	$\gamma = 0.2$; (μ_t, σ_t)	$\gamma = 0.05$; (μ_t, σ_t)
Data set 1	0.966, 0.003	0.947, 0.004	0.914, 0.007
Data set 2	0.980, 0.003	0.967, 0.005	0.942, 0.009
Data set 3	0.938, 0.003	0.906, 0.004	0.849, 0.010

This table shows the mean and the SD over the frames of the spatial correlation coefficient between the corrected images to their corresponding original ones. $N = 64$.

CONCLUSION

The performance of the proposed feature enhancement has been illustrated on two test images and on several echocardiographic sequences. An evaluation protocol has been proposed in the case of echocardiographic images and quantitative results have been presented. The proposed feature enhancement has one control parameter, γ; its influence on the results has also been studied. The consistency over time of the enhancement of the proposed approach has been demonstrated to ensure that no temporal artefacts are introduced. This is an important point, both for manual processing and assessment by a clinician, and for computer analysis of the sequence. The corrected images produced by the new method facilitate visual diagnosis by a clinician because the contrast between the heart wall and the cavity is enhanced. On the other hand, significant improvement in the results of the 2-D FA detection algorithm has been noted, in comparison with its application on the nonenhanced data.

The approach is original and gives good enhancement results on echocardiographic image sequences. The performance of the proposed approach is even more convincing if the processed data are displayed as a movie

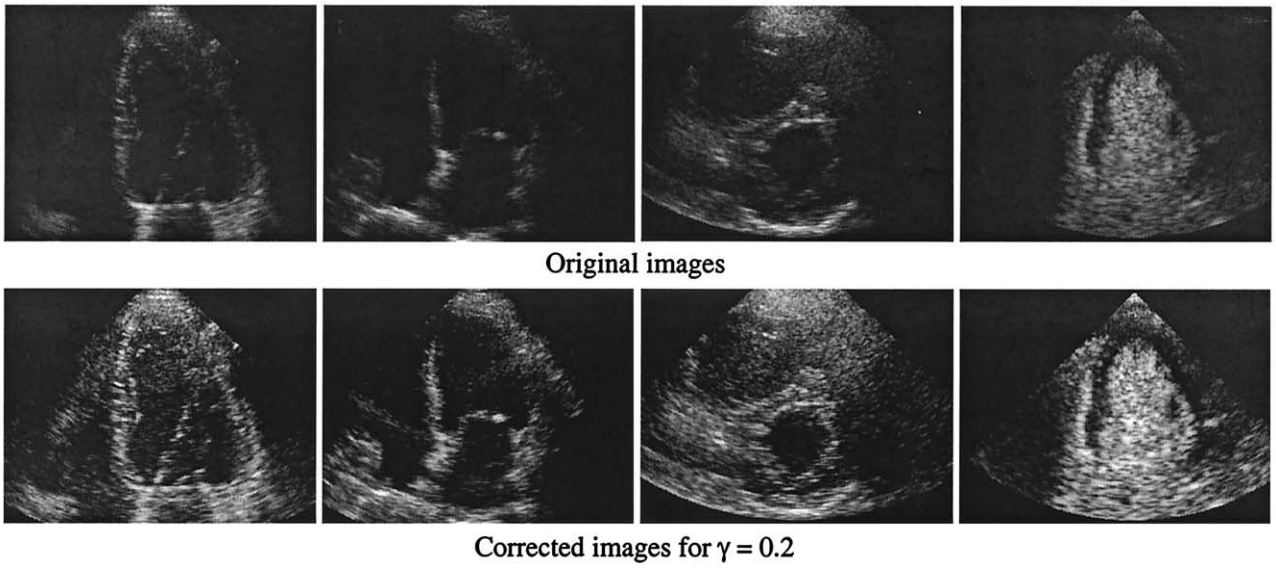

Fig. 10. Original and enhanced images of four different echocardiographic studies, $\gamma = 0.2$, $N = 64$. From left to right: B-mode long axis, B-mode four-chamber, B-mode short axis, CEPD long axis.

sequence.[1] The proposed approach is relatively fast in comparison with the existing algorithms in the literature (13 s/image 512 × 512 on a Pentium 400 MHz), but would require further speed-up for real-time processing and its use in a clinical environment (currently it takes 21 min to process one study of 100 frames).

Acknowledgements—The authors are grateful to Dr. M. M. Parada and Dr. J. Declerck, from Mirada Solutions Ltd, Oxford, for providing software used in part of this work. We also thank N. Clarke, from John Radcliffe Hospital, Oxford, and J. Bermejo, from Hospital General Universitario Gregorio Marañón, Madrid. This work was supported by the EC-funded ADEQUATE project (IST: 1999-10837).

REFERENCES

Ahmed MN, Yamany SM, Mohamed NA, Farag AA. A modified fuzzy c-means algorithm for MRI bias estimation and adaptive segmentation. Cambridge, UK: MICCAI, 1999:72–81.

Ashton EA, Parker KJ. Multiple resolution Bayesian segmentation of ultrasound images. Ultrason Imaging 1995;17:291–304.

Boukerroui D. Segmentation bayesienne d'images par une approche markovienne multiresolution. Application a l'imagerie échocardiographique et échographique 2D et 3D du sein. Ph.D. thesis. CREATIS, INSA 502. Villeurbanne Cedex (France), 2000.

Boukerroui D, Basset O, Baskurt A, Noble A. Segmentation of echocardiographic data. Multiresolution 2D and 3D algorithm based on gray level statistics. Cambridge, UK: MICCAI, 1999:516–524.

Cohen A, Daubechies I, Feauveau J-C. Biorthogonal bases of compactly supported wavelets. Commun Pure Appl Math 1992;45:485–560.

Daubechies I. Ten Lectures on Wavelets. CBMS-NSF Conference Series in Applied Mathematics, SIAM, Philadelphia, 1992:61.

Grimson WEL. From images to surfaces: Computational study of the human early visual system. Cambridge, MA: MIT Press, 1981.

Guillemaud R, Brady M. Estimating the bias field in MR images. IEEE Trans Med Imaging 1997;16:238–251.

Hughes DI, Duck FA. Automatic attenuation compensation for ultrasonic imaging. Ultrasound Med Biol 1997;23:651–664.

Kovesi P. Image feature from phase congruency. Videre: J Comp Vision Res 1999;1:1–26.

Leemput KV, Maes F, Vandermeulen D, Suetens P. Automated model-base bias field correction of MR images of the brain. IEEE Trans Med Imaging 1999;18:885–896.

Mallat S. A wavelet tour of signal processing. New York: Academic Press, 1999.

Morrone MC, Owens RA. Feature detection from local energy. Pattern Recogn Lett 1987;6:301–313.

Mulet-Parada M. Intensity independent feature extraction and tracking in echocardiographic sequences. Ph.D. thesis. Dept. of Eng. Science, Oxford University, 2000.

Mulet-Parada M, Noble JA. 2D+T acoustic boundary detection in echocardiography. Med Image Anal 2000;4:21–30.

Szeliski R. Fast surface interpolation using hierarchical basis functions. IEEE Trans Pattern Anal Machine Intell 1990;12:513–528.

Terzopoulos D. Multilevel computational process for visual surface reconstruction. Comput Vision Graphics Image Proc 1983;24:52–96.

Velthuizen RP, Heine JJ, Cantor AB, Lin H, Fletcher LM, Clarcke LP. Review and evaluation of MRI nonuniformity correction for brain tumor response measurements. Med Phys 1998;25:1655–1666.

Venkatesh S, Owens R. On the classification of image features. Pattern Recogn Lett 1990;11:339–349.

Wells WM, Grimson EL, Kikinis R, Jolesz FA. Adaptive segmentation of MRI data. IEEE Trans Med Imaging 1996;15:429–442.

Xiao G, Brady M, Noble JA, Zhang Y. Contrast enhancement and segmentation of ultrasound images—a statistical method. San Diego, USA: SPIE Med Imaging, Image Processing 2000; 3979: 1116–1125.

Yaou M-H, Chang W-T. Fast surface interpolation using multiresolution wavelet transform. IEEE Trans Pattern Anal Machine Intell 1994;16:673–688.

Zhang Y, Brady M, Smith S. Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. IEEE Trans Med Imaging 2001; 20:45–57.

[1] MPEG movies of the results are available at: http://www.robots.ox.ac.uk/~djamal/umb/umb.html