Available at www.ComputerScienceWeb.com
POWERED BY SCIENCE DIRECT®

Medical Image Analysis 7 (2003) 311–340

MEDICAL IMAGE ANALYSIS

www.elsevier.com/locate/media

Fusion of contrast-enhanced breast MR and mammographic imaging data

Christian P. Behrenbruch[a,b,*], Kostas Marias[a,b], Paul A. Armitage[a], Margaret Yam[a], Niall Moore[c], Ruth E. English[d], Jane Clarke[e], Michael Brady[a]

[a] *Medical Vision Laboratory (Robotics), Engineering Science, Oxford University, Parks Road, Oxford OX1 3PJ, UK*
[b] *Department of Surgery, Royal Free and University College Medical School, UCL, London NW3 2QG, UK*
[c] *Magnetic Resonance Imaging Centre, John Radcliffe Hospital, Headley Way, Oxford OX3 9DU, UK*
[d] *Breast Care Unit, Churchill Hospital, Oxford OX3 7LJ, UK*
[e] *Department of Surgery, Churchill Hospital, Oxford OX3 7LJ, UK*

Received 5 February 2001; received in revised form 25 February 2002; accepted 29 October 2002

Abstract

Increasing use is being made of Gd-DTPA contrast-enhanced magnetic resonance imaging for breast cancer assessment since it provides 3D functional information via pharmacokinetic interaction between contrast agent and tumour vascularity, and because it is applicable to women of all ages as well as patients with post-operative scarring. Contrast-enhanced MRI (CE-MRI) is complementary to conventional X-ray mammography, since it is a relatively low-resolution functional counterpart of a comparatively high-resolution 2D structural representation. However, despite the additional information provided by MRI, mammography is still an extremely important diagnostic imaging modality, particularly for several common conditions such as ductal carcinoma in situ (DCIS) where it has been shown that there is a strong correlation between microcalcification clusters and malignancy. Pathological indicators such as calcifications and fine spiculations are not visible in CE-MRI and therefore there is clinical and diagnostic value in fusing the high-resolution structural information available from mammography with the functional data acquired from MRI imaging. This paper presents a novel data fusion technique whereby medial–lateral oblique (MLO) and cranial–caudal (CC) mammograms (2D data) are registered to 3D contrast-enhanced MRI volumes. We utilise a combination of pharmacokinetic modelling, projection geometry, wavelet-based landmark detection and thin-plate spline non-rigid 'warping' to transform the coordinates of regions of interest (ROIs) from the 2D mammograms to the spatial reference frame of the contrast-enhanced MRI volume. Of key importance is the use of a flexible wavelet-based feature extraction technique that enables feature correspondences to be robustly determined between the very different image characteristics of X-ray mammography and MRI. An evaluation of the fusion framework is demonstrated with a series of clinical cases and a total of 14 patient examples.
© 2003 Elsevier B.V. All rights reserved.

1. Introduction

This paper introduces a system that has been developed to perform data fusion between 2D X-ray mammography and 3D contrast-enhanced MRI of the breast. The objective has been to develop a non-rigid registration and visualisation framework that makes use of both the high spatial resolution of mammography and the physiological (functional) information provided by contrast-enhanced MRI. This is particularly useful in situations where small opacities, microcalcifications and fine spiculations present in the X-ray image, cannot be correlated with CE-MRI voxel enhancement. Although the 2D–3D fusion process uses novel techniques, including a robust wavelet-based landmark detection method and 'pseudo X-ray' images generated from contrast-enhanced MRI, the actual non-rigid technique can be considered standard in the literature. In many respects, this article can be considered a 'clinical application' paper.

*Corresponding author.
E-mail address: cpb@robots.ox.ac.uk (C.P. Behrenbruch).

1361-8415/03/$ – see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/S1361-8415(03)00015-X

In general, there are several reasons for attempting to relate mammograms to 'functional' images, typified by contrast-enhanced breast MRI. These include:

(i) Complementary pathological indicators. For example, microcalcifications present in a mammogram may meaningfully correlate with an enhancing region in MRI, despite the fact that calcifications are *not* radiologically visible in MRI.
(ii) Patients having a pre-operative MRI scan often already have a mammogram and therefore there is the opportunity not only to provide the clinician with more specific information to support diagnosis but also to maintain diagnostic consistency across modalities.
(iii) Younger patients often have both an MRI scan and a mammogram due to the limited effectiveness of using mammography for assessing pre-menopausal women due to the X-ray attenuation characteristics of glandular tissue.
(iv) Post-operative tissue changes may be difficult to assess with mammography and are often better evaluated with MRI. This is largely due to the high X-ray attenuation of scar tissue.

The principal concept behind this work is the registration of a highly compressed 2D projective representation of the breast with an uncompressed volume-based acquisition. In the case of the MRI acquisition, the patient lies facedown in the scanner with the breasts pendulous and uncompressed in the breast coils. This is entirely different to the situation in mammography, where the patient normally remains upright with the breasts compressed (one at a time), between two plates. All of the patients in our study had both a medial–lateral oblique (MLO) and a cranial–caudal (CC) mammogram, as is required for breast cancer screening in the UK. It is worth noting that the angular separation between the mammogram view directions depends on the woman's size and shape, but typically varies between 45 and 60°. More importantly, as regards matching, the compression between the two mammograms differs significantly, often by as much as 1 cm.

In order to register the projective X-ray image with the MRI volume, we project the MRI volume in both the CC and the MLO direction. However, rather than simply use voxel intensity or approximate T1-correlated X-ray attenuation characteristics (Gilhuijs, 1995), we use an enhancement measure based on a two-compartment pharmacokinetic model of contrast uptake (Tofts and Kermode, 1991; Hayton et al., 1997; Brix et al., 1991; Buckley et al., 1994). This effectively produces a contrast-based projection that has the visual and structural characteristics of an 'enhanced' X-ray mammogram.

The registration technique is outlined in Fig. 1. It actually consists of two separate registration stages. The first 'partial registration' utilises a curvature measure to correlate boundary points along the film edge of the mammograms with the edge of the volume projection (Marias et al., 1999). This curvature measure depends on a good segmentation of the breast edge, which is achieved by an intensity-based search, mathematical morphology for

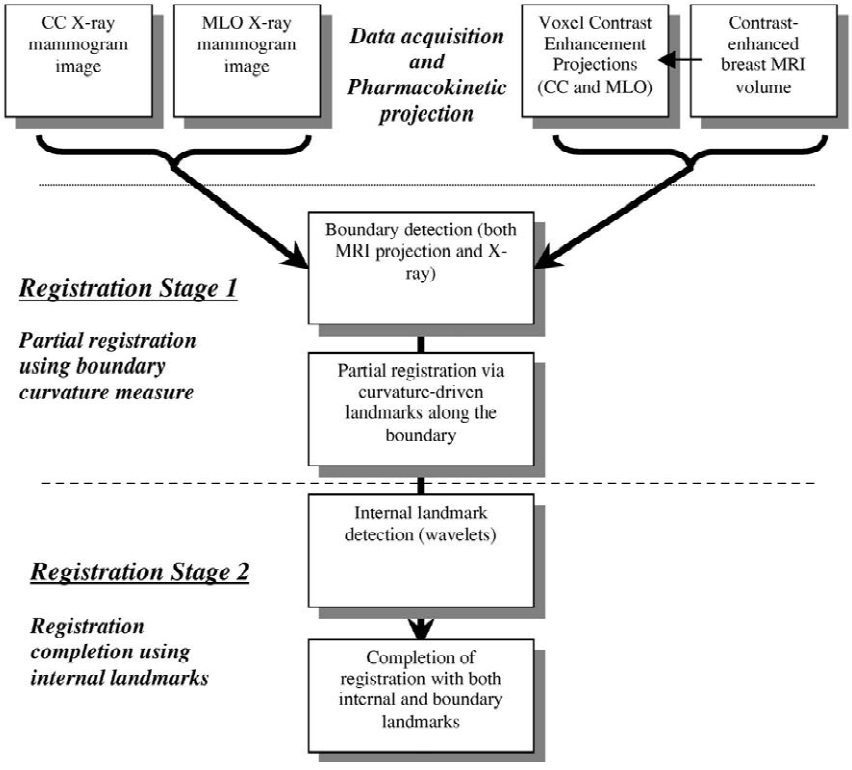


Fig. 1. Overview of X-ray/MRI data fusion illustrating the two stages of the registration process.

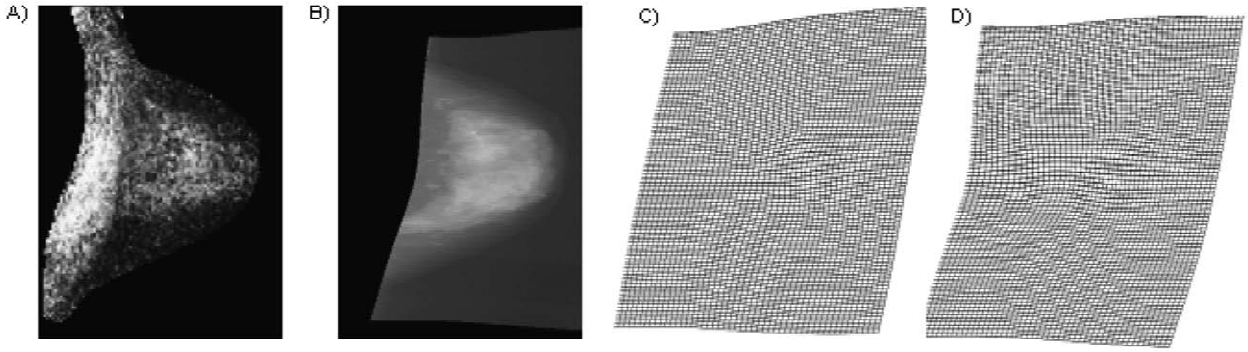


Fig. 2. (A) The functional projection of the MRI volume in the cranial–caudal plane. (B) The deformation of the cranial–caudal X-ray of the breast to the MRI contrast volume projection (*cropped*). (C,D) The deformation grid of the partially and fully registered images. The increased level of complexity in the deformation of (D) can be attributed to the internal landmarks used to complete the registration process.

smoothing and spline interpolation to produce a smoothed edge profile. The effect of this partial registration is to deform the extremities of the X-ray image to the boundary shape of the MRI contrast projection.

A second registration is performed based on the selection of internal landmarks using a wavelet-based feature detector. Internal landmarks in both the X-ray image and the MRI projection are matched using the following criteria:

- Scale localisation (i.e., landmarks of similar scale are matched).
- Orientation.
- Relative motion between the undeformed and the partially registered data from the first registration stage.
- Neighbourhood localisation in the partially registered images.

A wavelet-based technique has been used because features common to both the X-ray image and MRI projection may have a feature-space representation that has dramatically different scale characteristics. For example, a lesion visible in an X-ray mammogram may have fine spiculated details in addition to a main 'blob'-like characteristic, which is more representative of the resolution of the MR images. Additionally, there are many fine structural details in an X-ray image that do not have a feature correspondence with MRI and therefore a scale-sensitive method that has some 'adaptability' is important. Wavelet [packet] scale space with a best-basis function certainly provides a convenient and relatively easy to implement solution. The overall framework of the wavelet-based method used in this paper is not new, although the idea of a feature-driven registration algorithm which has a notion of 'feature scale' rather than merely taking a multi-scale image-driven approach is novel.

Registration completion involves using a combination of matched internal landmarks and breast boundary points. The registration technique adopted to date is thin plate spline warping (Bookstein, 1989; Evans et al., 1991), a variation of radial basis function (RBF) elastic deformation (Arad and Reisfeld, 1995). This landmark-driven approach is appropriate in applications where we do not necessarily have an intensity correlation between the MRI contrast projection and the X-ray data and therefore intensity-based techniques (Horn and Schunck, 1981; Thirion, 1998; Christensen et al., 1994) do not generally work well. This approach also enables us to perform the partial registration using only boundary points, which is particularly critical in the case where the breast has undergone significant involution[1] and there may be few landmarks or intensities suitable for controlling the internal deformation of the registration process. Typical registration (deformation) results of this two-stage registration process are shown in Fig. 2.

In the following sections, the main components of the registration process are discussed, namely the boundary curvature measure, internal landmark detection and the use of approximating thin-plate splines. Firstly, however, it is beneficial to discuss the notion of compression compensation, which is critical to our choice of registration framework.

2. Compression compensation

The underlying challenge in the breast 2D–3D data fusion process is compensation for the compression of the mammographically imaged breast. This is necessary as the MRI breast acquisition is not compressed (but is deformed/elongated by gravity) and therefore compensation is necessary to enable features visible in the mammograms to be spatially correlated with the MRI volume.

Previous work related to modelling breast compression has used a cross-section approach to uncompress the breast to enable point correspondence between mammogram

[1]The term 'involution' refers to the transformation of ductal or parenchymal breast tissue into fat. The incidence of involution increases sharply at the onset of the menopause. The transradiance of fat relative to ductal/parenchymal tissue is the reason screening is typically restricted to women over the age of 50, a heuristic approximation to the age of onset of the menopause.

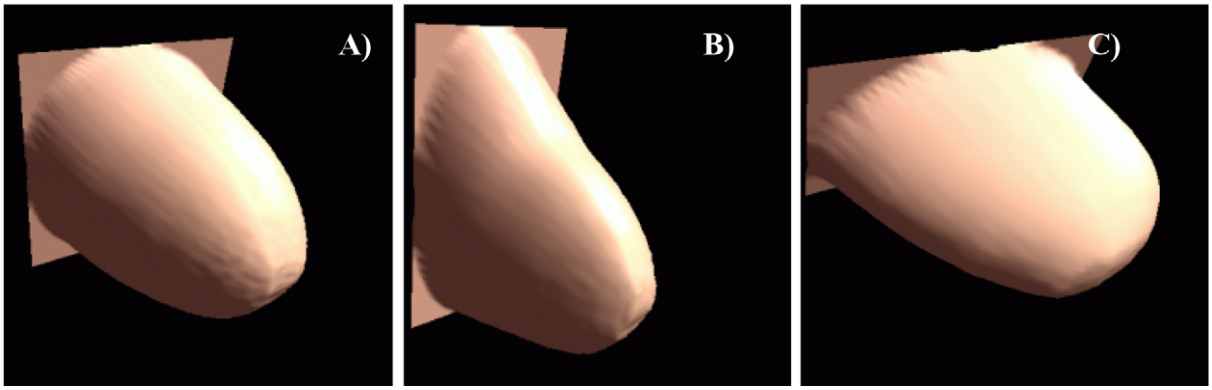


Fig. 3. (A) A 3D shape reconstruction based on the two uncompressed mammogram views (elliptical volume interpolation between views). Due to the fact that both the CC and MLO planes are compressed, the shape appears somewhat cylindrical. (B) A shape reconstruction based on the technique described in (Kita et al., 1998). (C) The shape of the breast based on the two mammograms registered to the MRI volume projections.

MLO (medial–lateral oblique) and CC (cranial–caudal) views (Yam et al., 1999; Kita et al., 1998). However, in our approach, the objective is to uncompress the mammograms into the reference frame of prone-acquired 3D MRI volumes and therefore a different strategy is required. The objective is to use the 'uncompression' of the mammogram images not only to facilitate a framework for point correspondence between mammograms, but also to enable pathology that is visible in the mammogram to be compared with MRI. As a consequence, the uncompressed shape is very different to that proposed by Kita et al. (1998) as the reconstructed shape resembles the prone-acquired MRI breast, rather than the idealised upright shape of the breast. This is illustrated in Fig. 3 where an uncompensated 3D shape reconstruction from an MLO and CC mammogram is compared with the model-based technique of Kita et al. and a reconstruction based on the mammograms registered to the MRI functional projections.

At first glance, Fig. 3(C), a two-view reconstruction using the registered X-rays, appears somewhat unusual. For comparison, Figs. 4 and 5 show the same registered and reconstructed views together with the segmented MRI projections. Not only is there good shape correspondence, but also it is possible to reconcile, albeit approximately, the location of the pathology between the two different modalities.

The registrations shown in Fig. 4 can then be used to reconstruct the approximate shape of the breast and compared against the MRI volume. Fig. 5 shows such a comparison, where the location of the tumour in the example used earlier is compared with the location of the tumour in the MRI volume segmented using a 3D im-

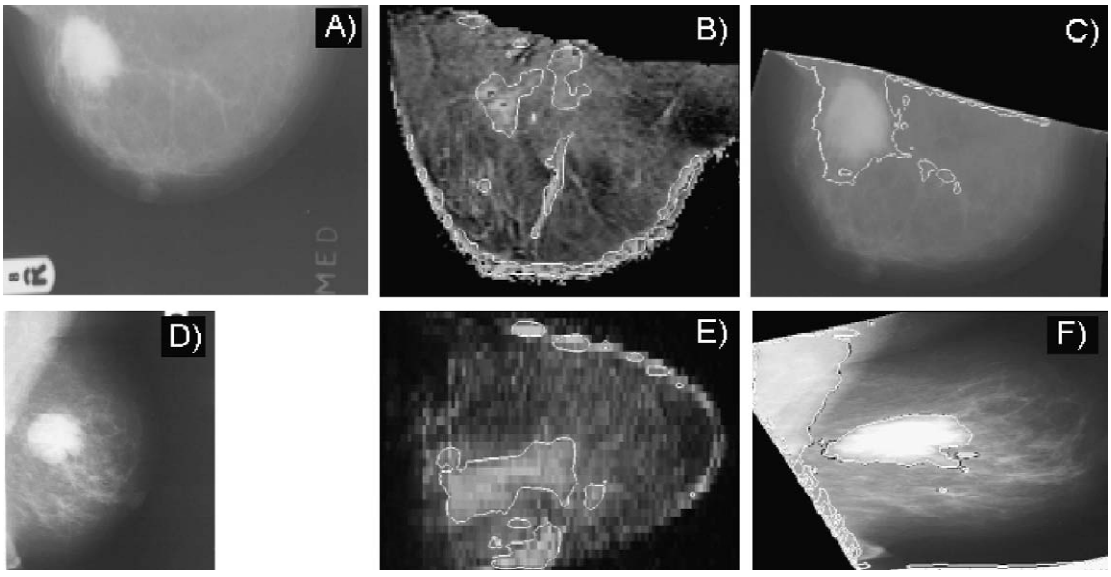


Fig. 4. (A,D) The original CC and MLO mammograms, respectively. (B,E) The segmented MRI functional projections. (C,F) The original mammograms registered to the MRI projections using the approach outlined in Fig. 1.

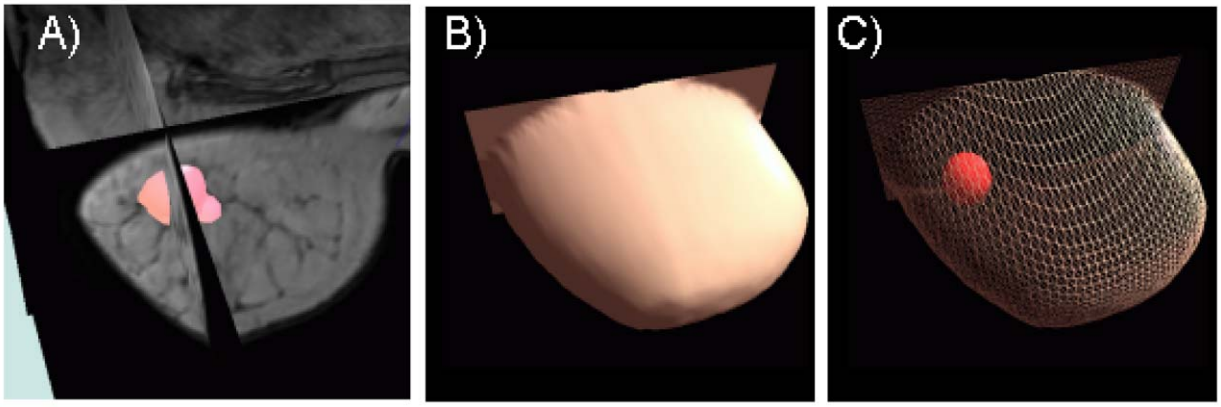


Fig. 5. (A) The location of the tumour in the (re-sliced) MRI volume. (B) The shape reconstruction displayed in Fig. 4(C). (C) The location of the mass in the 3D reconstruction based on an average of the CC and MLO view tumour centroid.

plementation of (Hayton et al., 1997). This example is of interest because it highlights a potential problem that must be addressed in the compression compensation approach, namely involution. If the breast is highly involuted, then it may not be possible to include more than the breast boundary (edge) in the model, which will most likely over-simplify the internal deformation of the breast.

3. Pharmacokinetic projection creation

3.1. Overview

The starting point for the fusion method introduced in this paper involves the simulation of mammograms from contrast-enhanced MRI, to enable comparison of the simulated results to actual mammograms. By geometrically projecting MRI voxel contrast-enhancement (pharmacokinetic) information (Fig. 6) to a cranial–caudal and/or medial–lateral (oblique) view, and registering the mammograms to the resulting projection, it is possible to compare the functional structure of tissues with the high-resolution structure of X-rays. The registration of X-rays to the reference frame of MRI enables not only correspondence between mammogram views, but also approximates correspondence with imaged tissue using MRI.

3.2. Projection formation

The main idea behind the projection step is to create an image from an MRI breast volume that is geometrically similar to a mammographic projection, but which represents the breast in an uncompressed shape. Rather than simply using an intensity-based projection or some kind of transfer function that maps MRI T1 values to approximate tissue X-ray attenuation characteristics, we project the voxel enhancement characteristics using a pharmacokinetic model of gadolinium uptake. This concept is illustrated below, where low enhancement tissues are given a low level of intensity in the ‘pseudo X-ray’ projection.

The projected intensity from high-enhancement voxels is given a correspondingly high value. To map the projection

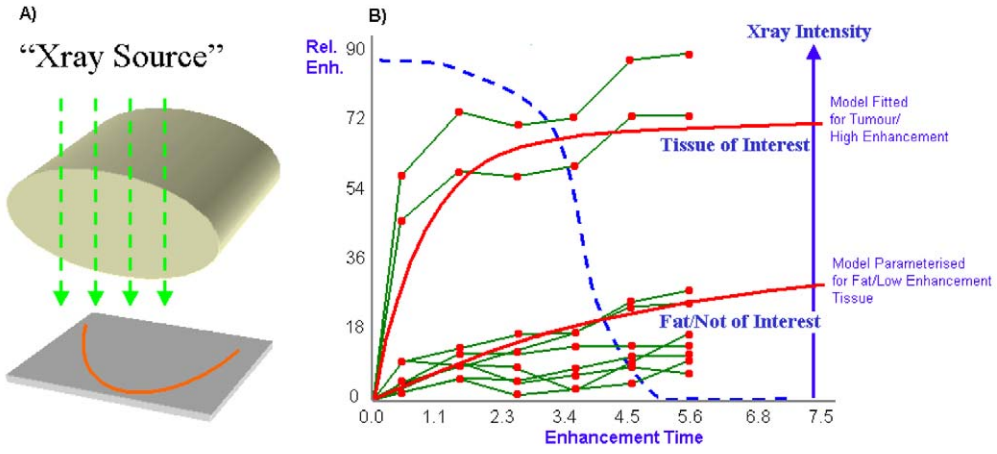


Fig. 6. (A) The conceptual geometry behind the projection process (for a cranial–caudal view), where the MRI voxel intensities are mapped into a transformed intensity plane. (B) Some typical voxel enhancement graphs with the fitted pharmacokinetic model. The dashed line shows the effective result of integrating the fitted model to produce a projection intensity that is biased to voxels with high enhancement.

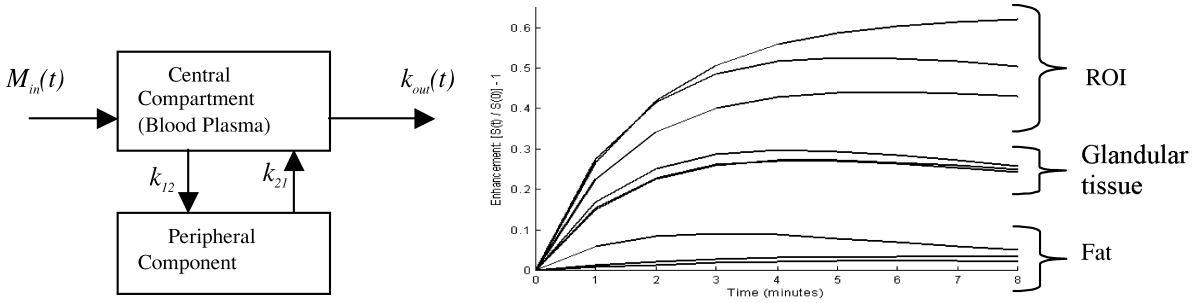


Fig. 7. A two-compartment pharmacokinetic model with typical contrast curves for fat, parenchymal (glandular) tissue and enhancing regions of interest. M_{in} is the mass of contrast injected into the blood stream with respect to time. k_{12} and k_{21} are inter-compartment exchange rates and k_{out} is the contrast excretion rate.

intensities, we integrate the Hayton–Brady pharmacokinetic model of contrast uptake (Hayton et al., 1997) for each voxel with respect to enhancement time. The various components of this pharmacokinetic model are depicted in Fig. 7.

Eq. (1) shows the derived bolus model for Gd-DTPA uptake, which is fitted to the contrast data (via a Levenberg–Marquardt algorithm) to produce a difference of exponential parameterisation (see Fig. 7(b) for example curve fittings).

$$\mathrm{Conc}(t) = \frac{A}{a-b}\left(\exp^{-bt} - \exp^{-at}\right). \quad (1)$$

The parameters A, a and b characterise the exchange rates k_{12}, k_{21} and k_{out} (Fig. 7(a)) although for the purposes of controlling the projection, we are primarily interested in the rate and amplitude of the voxel contrast enhancement profile.

3.3. Projection functional characteristics

The result of the pharmacokinetic projection is a functional projection of the MRI volume that has a number of interesting characteristics. These are:

- Fatty tissue receives a low integrated projection enhancement.
- Parenchymal (glandular) tissue is given an enhancement level that spatially corresponds well with the intensities that are observed in a mammogram image.
- Rapidly enhancing regions of interest (ROIs) are given extremely high relative contrast.

These characteristics are demonstrated in the two cranial–caudal examples shown in Fig. 8. In the first pair, the patient has no pathology (except for a few coarse calcifications); however, the glandular tissue is strongly pronounced in the mammogram due to ductal ectasia. These ‘exploded’ ducts are clearly visible in the contrast-enhanced MRI projection as a correspondingly high level of contrast with respect to the fatty background. In the second example, the patient has a large, poorly differentiated tumour (outlined) that has a similar intensity to glandular tissue in the X-ray, but which clearly stands out as a strongly enhancing region in the MRI projection. It is worth noting that the comparative resolution between the projections is quite different due to the spatial resolution of the MRI scan (particularly the number of slices). The projected intensity distribution also varies, as it is dependent on the degree of involution in the breast.

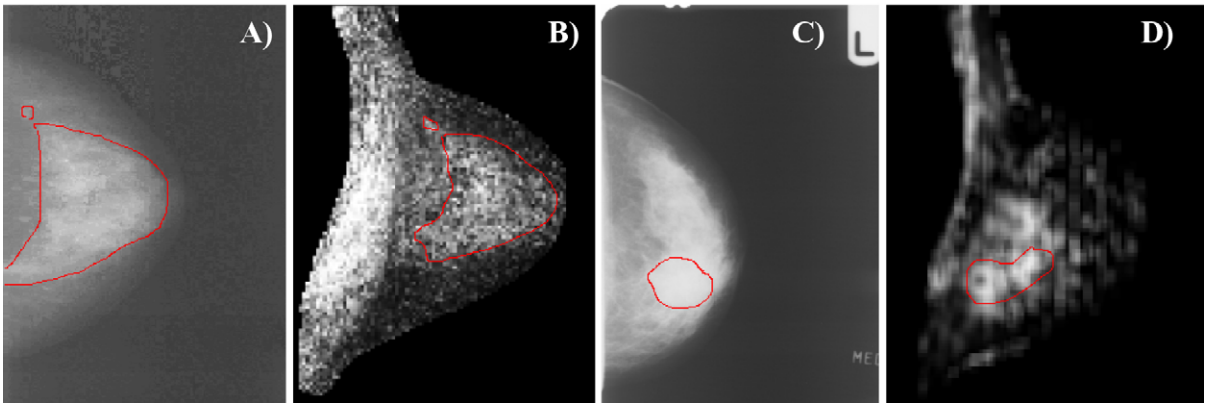


Fig. 8. Two examples comparing the CC mammogram to corresponding contrast-enhanced MRI projections. The first example (A,B) shows consistent shape characteristics of moderately enhancing parenchymal tissue, due to ductal ectasia. The second pair (C,D) compares the contrast of a large tumour in the X-ray image with the corresponding functional projection.

3.4. Alternative pharmacokinetic approaches

Recently, we have considered alternative pharmacokinetic models (Armitage et al., 2001). The most simplistic of these is a statistical comparison of the integral of voxel enhancement. Although this provides reasonable tissue differentiation, the lack of a physiologically meaningful model to constrain the analysis renders the model highly susceptible to noise, particularly for low-enhancing fatty regions.

The model that has, to date, provided satisfactory results is the Hayton–Brady model. As discussed previously, this method is based on a relatively conventional two-compartment exchange model. More complex models exist that include additional exchange components, notably that proposed by Tofts and Kermode (1991). It should be noted, however, that low temporal resolution of the enhancement acquisition (7–9 time points, both breasts imaged at approximately 1-min intervals) does not support system identification with a more complex model representation. Of course, improved temporal resolution of the MRI acquisition sequence resulting from a small field of view (FOV) or a single breast acquisition would make complex models more applicable. In fact, such sequences are already widely in use. However, current clinical judgement accords greater emphasis to spatial rather than temporal resolution and both breasts are always imaged in case the disease is multi-focal.

3.5. Motion correction prior to projection

There are two additional issues that require consideration in order to accurately compare the mammograms to their corresponding MRI functional projections. The first is motion correction. In order to evaluate the voxel enhancement with respect to time, it is necessary to ensure that tissue motion is limited. Since the CE-MRI acquisition period (including pre-contrast agent acquisitions) takes around 10 min, it is unrealistic to expect that the patient will remain completely motionless. This is a particular problem as symptomatic women are often naturally tense when imaging begins but relax several minutes into the procedure, resulting not only in general body (and hence breast) movement but internal soft-tissue deformation due to pectoral muscle flexation. A number of research groups have developed approaches to this problem including techniques utilising (non-rigid) optical flow (Hayton et al., 1997) and tensor splines (affine plus local deformation) as a motion correction framework (Reuckert et al., 1999). We currently utilise a 3D variant of the approach developed by Hayton et al., which uses mutual information as a voxel similarity measure combined with a novel Bayesian regularisation method to perform optical flow registration.

4. Breast boundary feature detection

The detection of suitable landmarks for boundary deformation between the X-ray mammogram and the MRI volume projection first involves segmenting the breast edge. To achieve this, the X-ray images are binary thresholded at the first ‘valley’ of the intensity histogram using a simple gradient/slope detector as shown in Fig. 9(a). This approach also works for the MRI projection although it is essentially unnecessary as background (air) voxels are easy to threshold prior to projection. Following segmentation of the breast area we generate an eight-connected component outline by using mathematical morphology (Sallam and Bowyer, 1996) (closing followed by dilation and subtraction from the original binary), shown in Fig. 9(c).

Although this method results in an approximation of the boundary, it has the virtue of being very fast and the resulting outline is smooth (due to morphology) while preserving the anatomical features of the original image (for example, the location of the nipple). Once the breast

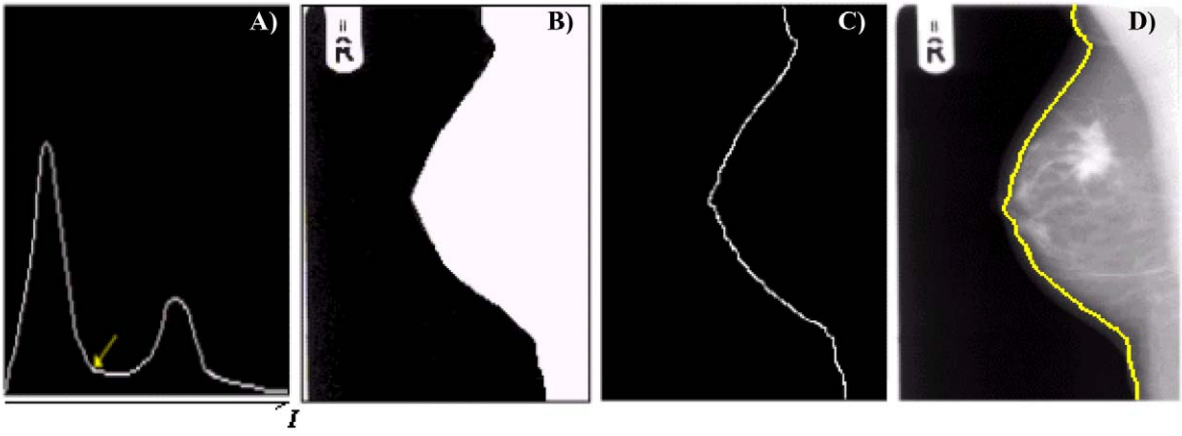


Fig. 9. The steps to generate breast boundary segmentation involve (A) bimodal histogram thresholding, (B) binary image creation and (C) morphology to produce the breast outline. (D) The outline over the breast superimposed on the original image.

outline has been detected it is sampled and approximated with a cubic spline $g^i(s)$ that is parameterised along the edge (via a simple graph search) for every three points of the breast boundary. This parameterisation enables the boundary curvature $K(s)$ to be calculated for each point along the spline:

$$K(s) = \frac{x_{ss} \cdot y_s - x_s \cdot y_{ss}}{x_s^2 + y_s^2}. \quad (2)$$

Previous work with mammography registration has focused predominantly on aligning the boundary of the breast in each mammogram (Marias et al., 1999; Sallam and Bowyer, 1996; Wirth and Choi, 2000). In our application, we automatically calculate the location of the nipple and the points between the breast, the rib cage and the axilla (in both the mammogram and the projection) on the basis of separating positive and negative curvature portions of the spline parameterisation and subsequently locate curvature peaks as shown in Fig. 10. Separation of the positive and negative curvature components from Eq. (2) is very important as it enables the detection of consistent landmarks near the rib and the axilla.

These curvature-selected points form the basis of the first step in the registration process, which effectively ensures that the boundary of the X-ray image is warped onto the MRI volume projection. A non-rigid transformation is made possible by interval subsampling at regular intervals between the curvature-detected landmarks (5–10 points per segment, depending on the size of the breast), in order to generate a large number of points for 'warping' the boundary between the X-ray mammograms and the 'pseudo X-ray' MRI projections. Our boundary detection technique has been reliably applied to many mammograms (Marias et al., 2000) and is sensible to use in this fusion approach due to the outline shape consistency between the MRI projection and the true mammogram. However, we have found that this is not an adequate correction for breast compression, so that internal landmarks are required to complete the compensation for complex internal soft tissue deformation.

5. Partial registration from the breast boundary

5.1. Overview

Radial basis functions (RBFs) are used for the elastic image deformation in this registration scheme. RBFs are used in two contexts—firstly, in an interpolation framework for the boundary conditions where the regularisation requirements are simple. Later when internal landmarks are included, information about the spatial characteristics of the deformation points is used to implement a more sophisticated regularisation that is based on an approximation method. We start with the interpolation case as it is more general to defining a RBF.

5.2. Interpolative radial basis functions

In RBF interpolation, a set of n landmarks (p_i, q_i) is used to define a transformation function $u: R^2 \rightarrow R^2$ where $p_i = (x_{i1}, y_{i1})$ are the landmarks in the first image and $q_i = (x_{i2}, y_{i2})$ are the landmarks in the second. An interpo-

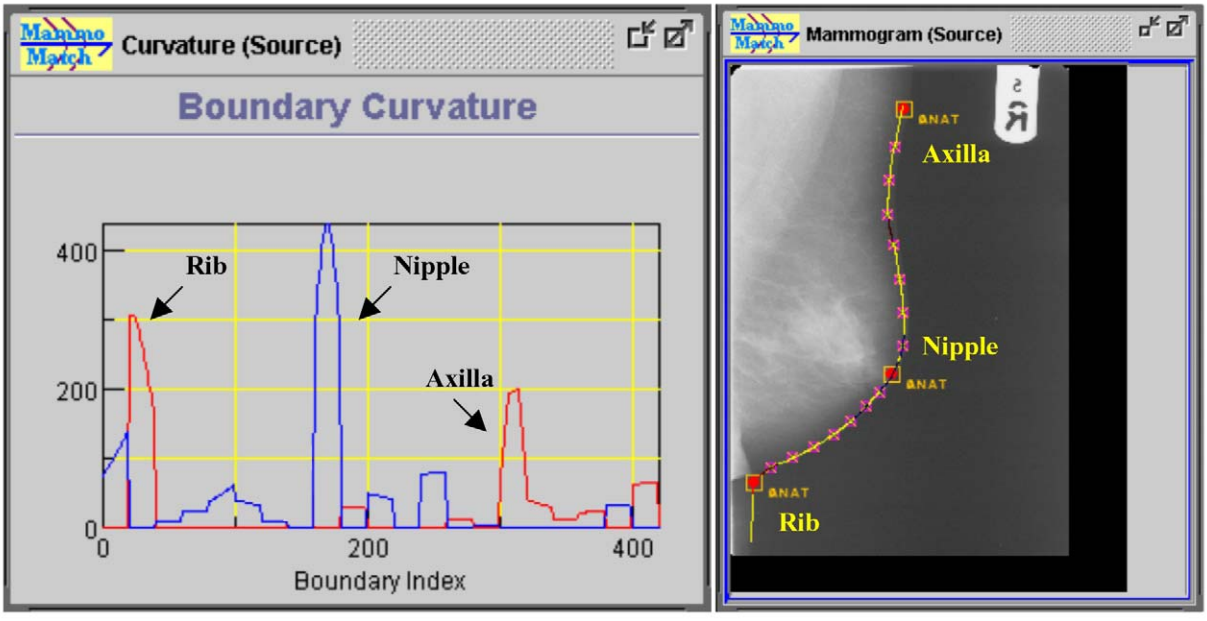


Fig. 10. The curvature profile along the breast boundary $K(s)$, consisting of the two peaks near both ends of the boundary and the peak corresponding to the nipple in the middle. The boundary is then sub-sampled at intervals between the rib-nipple and nipple-axilla segment to produce enough points to 'warp' the boundaries.

lation transformation function $u(x)$ based on point landmarks must fulfil the following constraint:

$$u(p_i) = q_i, \quad i = 1, \ldots, n, \tag{3}$$

and minimise a functional $J(f)$ for R^2.

$$J_2^2(f) = \int\int_{R^2} \left\{ \left(\frac{\partial^2 f}{\partial x^2} \right)^2 + \left(\frac{\partial^2 f}{\partial y^2} \right)^2 + 2 \left(\frac{\partial^2 f}{\partial x \cdot \partial y} \right)^2 \right\} \mathrm{d}x\, \mathrm{d}y, \tag{4}$$

where this functional is also known as the bending energy of the deformation. In this case we are solving the biharmonic equation $\Delta^2 f = 0$ whereby local deformation conforms to a function $f(x,y)$ for minimisation of the bending energy. For the interpolation case, each coordinate is calculated independently (i.e., the interpolation $u_k{:}R^d \to R$ is solved for each coordinate $k = 1 \ldots d$, in this case $d = 2$, with the corresponding constraints $u_k(p_i) = q_i,k$). In the general RBF approach, the interpolation function $u(x)$ consists of two parts:

$$u(x) = \phi_s(x) + R_s(x), \tag{5}$$

where $\phi_s(x)$ is a sum of polynomials up to degree p and $R_s(x)$ consists of a sum of RBFs:

$$\phi_s(x) = \sum_{j=1}^{M} \beta_j \phi_j(x), \tag{6}$$

$$R_s(x) = \sum_{i=1}^{N} \alpha_i R(\|x - p_i\|). \tag{7}$$

The $\phi_j(x)$ are a separable basis of M functions for all polynomials up to degree p. $R(r) = R(\|r\|)$ is a distance function (depending only on $r \geqslant 0$ from the origin) that describes the local deformation (interpolation) characteristics of points around the RBF landmarks. In Eq. (7), $\|x - p_i\| = \|r\|$ is the Euclidian distance from x to p_i and α_I, β_j are coefficients, with the RBFs $R(\|x - p_i\|)$ centred on the N landmarks p_i. Combining (6) and (7) and using the deformation constraint:

$$\sum_{i=1}^{n} \alpha_i \phi_i(p_i) = 0, \quad j = 1, \ldots, M, \tag{8}$$

yields a system of linear equations for the coefficients α, β:

$$\begin{pmatrix} K & P \\ P^{\mathrm{T}} & 0 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} q_k \\ 0 \end{pmatrix}, \tag{9}$$

where K is the $n \times n$ matrix $K_{ij} = R(\|p_i - p_j\|)$ and P is the $N \times M$ matrix $P_{ij} = \phi_j(p_i)$. The $q_k = (q_{k,1}, \ldots, q_{k,n})^{\mathrm{T}}$ are the set of target landmarks.

There are a wide variety of radial basis functions used in the literature for elastic deformation of images. These include thin-plate splines (R_{TPS}) (Bookstein, 1989; Rohr et al., 1996), multiquadrics (R_{M}) (Hardy, 1981; Little et al., 1997), inverse multi-quadrics (R_{IM}) (Ruprecht and Müller, 1993) and Gaussian RBFs (R_{G}) (Arad and Reisfeld, 1995):

$$R_{\mathrm{TPS}}(r) = \begin{cases} r^{4-d} \ln r^2, & 4 - d \in 2N, \\ r^{4-d}, & \text{otherwise}, \end{cases} \tag{10}$$

$$R_{\mathrm{M}}(r) = (r^2 + c^2)^{\mu}, \quad \mu \in R_+, \tag{11}$$

$$R_{\mathrm{IM}}(r) = (r^2 + c^2)^{-\mu}, \quad \mu \in R_+, \tag{12}$$

$$R_{\mathrm{G}}(r) = \mathrm{e}^{-r^2/2\sigma^2}. \tag{13}$$

In our work to date, we have used the thin-plate spline radial basis function R_{TPS} (10), since the effect of deformation is both intuitively reasonable and (qualitatively) produces results that appear clinically consistent. Control of regularisation is extended in the second stage of the registration process, but it is conceivable that with proper parameterisation of the distance functions, it would be possible to use other RBFs (for example, a Gaussian) in this application.

6. Wavelet-based internal landmark detection

6.1. Overview

To complete the registration process, corresponding internal landmarks are detected in the registration pair and, along with various matching conditions (discussed in the next section), are used to compensate for the complex internal deformation of the breast. We cannot rely on a field-based (intensity-based) registration technique as the deformation may be quite large and the local textural detail of features is different in the X-ray and corresponding CE-MRI projection. Therefore it is necessary to implement a robust feature detection algorithm, which in turn produces landmarks that can be matched and used as the basis for non-rigid registration refinement.

Various spatial filtering techniques based on localised Fourier operators and trigonometric transforms have been used for image feature detection (D'Astous and Jernigan, 1984; Haralick et al., 1973; Turner, 1986; Bovik et al., 1990). The inherent limitation of these approaches lies in the need to define a large number of Fourier windows (in a multi-resolution sense) in order to produce a maximally sensitive feature detector. In order for these techniques to work robustly, there is a need to be able to adapt the spatial localisation properties of filters as meaningful descriptors with respect to image content. Most windowed Fourier/Gabor techniques can do this only to a limited extent.

An alternative to windowed Fourier techniques is the discrete wavelet transform (DWT), a dyadic filtering implementation that provides a true multiscale foundation for image analysis (Mallat, 1998). A useful variant of the

DWT is 'discrete wavelet packet analysis' (DWPA) (Coifman et al., 1991, 1992) which is essentially a reconfiguration of the basis subspaces. This results in particular linear combinations (superposition) of wavelets, producing a library of wavelet projections (called 'packets') with which to analyse a function. Additionally, DWPA forms subspaces that retain the orthogonality, smoothness and localisation properties of the dyadic DWT, but with better spatial localisation characteristics.

The improved spatial characteristics are demonstrated in an example in Fig. 11, where an (approximate) impulse function is decomposed using both wavelets and wavelet packets. When viewed as a tiling of information cells across the space-frequency plane, the DWT analysis correctly localises the peak of the function in the high frequency elements, but is forced to include poorly localised, low frequency components as well. Although the analysis equivalent for the wavelet packet case does not contain all spatial-frequency tilings, it demonstrates that it is possible to reconstruct the signal using almost a single wavelet packet. For this reason, we have used DWPA as the foundation of our feature detection, particularly as we are interested in highly localised features such as calcifications and spicules are good quality, consistent landmarks for registration.

Using H and G to represent the equivalent high and low pass filter components of a set of conjugate quadrature filters, the following recursive set of functions can be defined:

$$\psi_0 \stackrel{\text{def}}{=} H\psi_0; \quad \int_R \psi_0(t)\,\mathrm{d}t = 1, \tag{14}$$

$$\psi_{2n} \stackrel{\text{def}}{=} H\psi_n; \quad \psi_{2n}(t) = \sqrt{2}\sum_{j\in Z} h(j)\psi_n(2t-j), \tag{15}$$

$$\psi_{2n+1} \stackrel{\text{def}}{=} G\psi_n; \quad \psi_{2n+1}(t) = \sqrt{2}\sum_{j\in Z} g(j)\psi_n(2t-j). \tag{16}$$

The function ψ_0 (the 'root') is uniquely defined and is identified with wavelet generation function ϕ in (Mallat, 1989). The function ψ_1 is the 'mother wavelet' associated by convolution with H and G. The descendent functions (tree 'leaves') ψ_n are analogous to iterated forms of the mother wavelet as in the DWT; however, in this case, the evolution of ψ_n is over the complete tree. The collection of these functions for $n = 0,1,\ldots,m$ makes up the wavelet packets associated with a particular choice of quadrature filters. It is also apparent that in this way, wavelet packet analysis is a generalisation that includes the DWT decomposition

The most straightforward way to use wavelet packets is in a fixed scale context. In this way, they are somewhat analogous to analysing a function or sequence with a DWT at a fixed scale or level, albeit with a different set of filter projections. In our case, we use wavelet packets in a multiscale context, which effectively means that the approximation can be performed using projections anywhere in the decomposed basis. For further details of the concepts of wavelet analysis and scale-space techniques, we refer to (Witkin, 1983; Burt and Adelson, 1983; Daubechies et al., 1986; Stein, 1979; Marr, 1982; Grossman and Morlet, 1985; Mallat, 1998) and for a very useful overview of signal processing with wavelet packets (as well as a good treatment of the software implementation issues) (Wickerhauser, 1994).

6.2. Best basis selection

The multiscale wavelet packet representation provides a rich set of filtered projections within a wavelet basis at different scales. Using wavelet packets for feature detection involves ordering these subspaces by information content to produce a scale-space 'stack' for analysis. This approach is called 'best basis analysis' (Coifman and Wickerhauser, 1992). Before constructing such a scale space, it is necessary to define the notion of an information cost function that is used to rank each representation in terms of the contribution to approximating a sequence. To

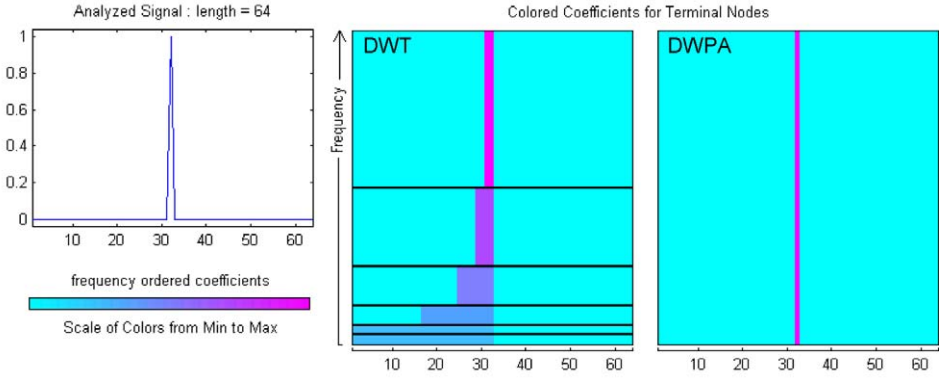


Fig. 11. A comparison between DWT and DWPA representation of an (approximate) impulse function. The DWT case (left) correctly determines the location of the peak for the high-frequency components but also includes low frequency components with poor spatial localisation. The DWPA case (right) demonstrates much better spatial localisation and produces a more optimal representation of the impulse.

do this, we use a real-valued functional M, satisfying the condition

$$M(u) = \sum_{k \in Z} \mu(|u(k)|); \quad \mu(0) = 0. \tag{17}$$

Here μ is a real-valued function defined on $[0,\infty)$. If we assume that $\Sigma_k \mu(|u(k)|)$ converges absolutely, then M will be invariant to rearrangement of the representation of the sequence u. We can then define a functional M_x on the set of bases B for each $x \in X$:

$$M_x: B \to R; \quad B \to M(B^*x), \tag{18}$$

which is referred to as the *M-information cost* of a sequence x in the basis B. Another way of looking at this, is to evaluate a transformation of $f(x)$ to $g(x)$ via an operator (in this case convolution–decimation) as a distance measure with respect to a transformation parameter λ.

$$\|f - g\|^2_{L_2(I)} + \lambda \|g\|, \tag{19}$$

evaluated as

$$\|f - g\|^2_{L_2(I)} \triangleq \left(\int_I |f(x) - g(x)|^2 \, dx \right)^{1/2}. \tag{20}$$

This enables each approximated component of the wavelet packet decomposition to be compared and subsequently ranked by information content. There are many different kinds of information cost functionals that can be used. For example, a simple threshold ε based on coefficients in a wavelet subspace could be used to order the contribution of each packet to the complete reconstruction of a sequence, as shown in Eq. (21).

$$\mu(\omega) = \begin{cases} |\omega|, & \text{if } |\omega| \geqslant \varepsilon, \\ 0, & \text{if } |\omega| < \varepsilon. \end{cases} \tag{21}$$

Another useful approach is the entropy of a representation, such that a cost function $u = \{u(k)\}$ is evaluated by

$$H(u) = \sum_k p(k) \log \frac{1}{p(k)}, \tag{22}$$

where $p(k) = |u(k)|^2 / \|u\|^2$ is the normalised energy of the kth element of the sequence. This functional can be evaluated in a minimisation context and used to rank a set of decompositions.

One of the difficulties with using conventional entropy measurements is defining a probability density function p in a way that is appropriate to each representation. Practically, this means selecting a histogram sampling density that meaningfully captures the statistical characteristics of the sequence. Clearly, in a multi-scale context where this distribution may vary greatly (and with different sampling resolutions by virtue of decimation), this is not a trivial task. To circumvent this problem, we have chosen to use an approximation to entropy—*Approximate Entropy* (ApEn), which is a useful measure of signal regularity (Pincus, 1992, 1997). The idea is that wavelet packet decompositions that smoothly approximate the original sequence will have a higher degree of regularity, whilst equivalent high-pass filtered components that reflect the signal noise and will be more fluctuating in nature, exhibiting greater irregularity. ApEn is defined as follows:

$$\Phi^m(r) = \frac{1}{N - m + 1} \sum_{i=1}^{N-m+1} \log C_i^m(r), \tag{23}$$

where given a positive integer N and nonnegative integer m, with $m \leqslant N$, a positive real number r, and a sequence of real numbers $u := (u(1),u(2), \ldots ,u(N))$. The distance between two sequences $x(i)$ and $x(j)$, where $x(i) = (u(i),u(i+1), \ldots ,u(i+m-1))$ is defined by $d(x(i),x(j)) = \max_{p=1,2,\ldots,m}(|u(i+p-1) - u(j+p-1)|)$. $C_i^m(r)$ is the number of $j \leqslant N - m + 1$ such that $d(x(i),x(j)) \leqslant r$, normalised with respect to $N - m + 1$. Thus, ApEn$(m,r,N)(u)$ may be interpreted as a measure of the frequency at which number sequences of length m occur compared with sequences of length $m + 1$. High values of the ApEn measure imply a highly random nature, whereas low values suggest a relatively smooth sequence.

Therefore, if we apply the ApEn measure to each wavelet packet decomposition of a sequence, we can order each projection by the contribution it makes to completely reconstructing the sequence. Starting from the lowest approximate entropy packet decompositions and progressively adding detail (higher ApEn), one can construct a scale-space using multi-scale wavelet packets, despite the fact that different scales (depending on whether they are equivalent low-pass or high-pass filtering combinations) may be ranked counter-intuitively by virtue of the level of decomposition. An example of such a scale-space is shown in Fig. 12, where a small portion of an X-ray image is decomposed into a wavelet packet scale-space stack (2D sequences are discussed more thoroughly in the next section). It is worth noting that the resulting scale space appears similar to many other ‘scaleogram’ techniques, including phase-based methods such as (Kovesi, 1996). The difference with our method lies in the fact that best-basis selection pre-orders the scale decomposition in a fashion that is more oriented to the feature content of the image (by virtue of using morphologically relevant wavelets). The result is a feature-driven scale space that dramatically improves the ease of defining salient regions.

6.3. Choice of wavelet packet

Because of the superposition of wavelets, wavelet packets provide a large range of multi-scale representation with which to create a scale space. For this work, we have chosen a well-known class of existing wavelet bases, called ‘Coiflets’, which have a number of useful properties that make them ideal for detecting regions of interest in medical images. In particular, we wanted a wavelet basis

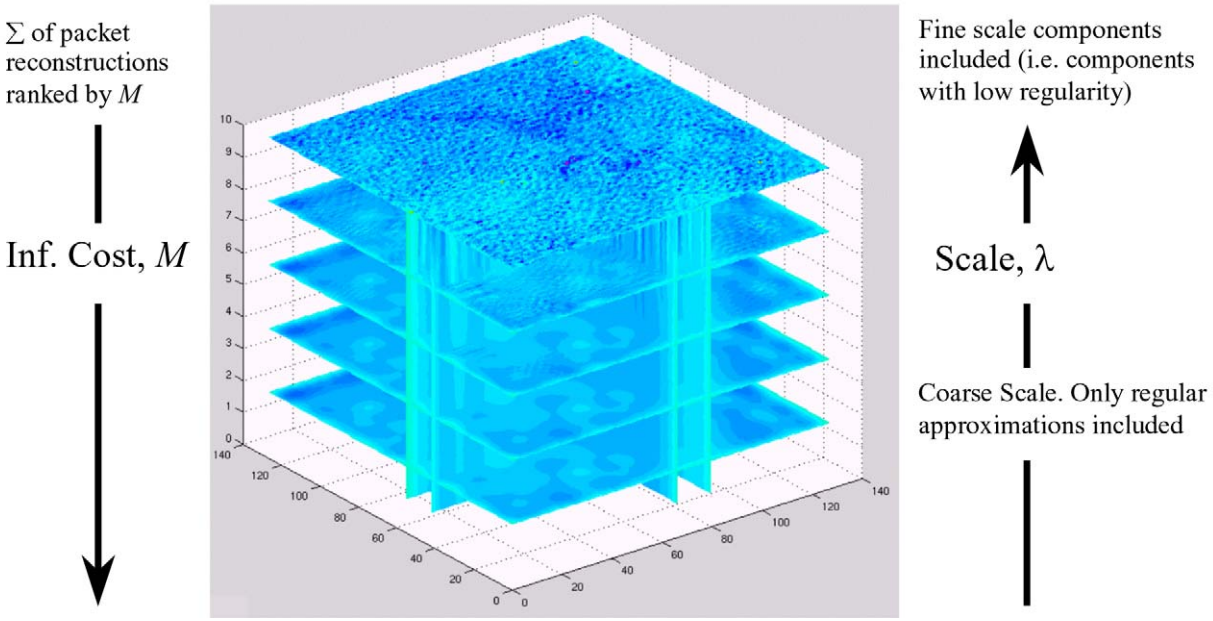


Fig. 12. An example scale-space illustrating cumulative addition of image detail (via packet reconstructions) using an information cost function. The bottom of the scale-space contains a quite smoothed approximation to the image, whilst the top of the stack is a perfectly reconstructed image (i.e., all packets have been used).

that would be sensitive to sharp, spatially localised features such as tumour spiculations, microcalcifications (in X-ray images) and small regions of focal enhancement in MRI (such as for ductal carcinoma). Coiflets have the following useful properties:

- They have a high number of vanishing moments for a given support length (i.e., good frequency localisation as well as spatial localisation).
- They are almost symmetrical.
- They have compact support.
- They are orthogonal (indeed a bi-orthogonal form is also available).

Low-order Coiflet wavelet packets are useful for approximating 'bumps' without over-smoothing, hence they have good edge-preservation characteristics. Higher-order packets have broad frequency characteristics making them useful for noise approximation. These properties make Coiflets very good for generating the kind of scale-space shown in Fig. 12, as a good spread of detail is ascribed to both the smoothly approximated region of the scale-space and the high-frequency (noise) components. In the case of features that have a broad Fourier localisation (for example, sharp changes in intensity due to edges or highly localised features), the Coiflet scale-space provides good discrimination from smoother components. Some example Coiflet wavelet packets are shown in Fig. 13 (generation and scaling functions shown, as well as example higher-order packets). For more information on the generation and properties of Coiflets we refer to (Daubechies, 1992). Additionally we are considering the development of cus-

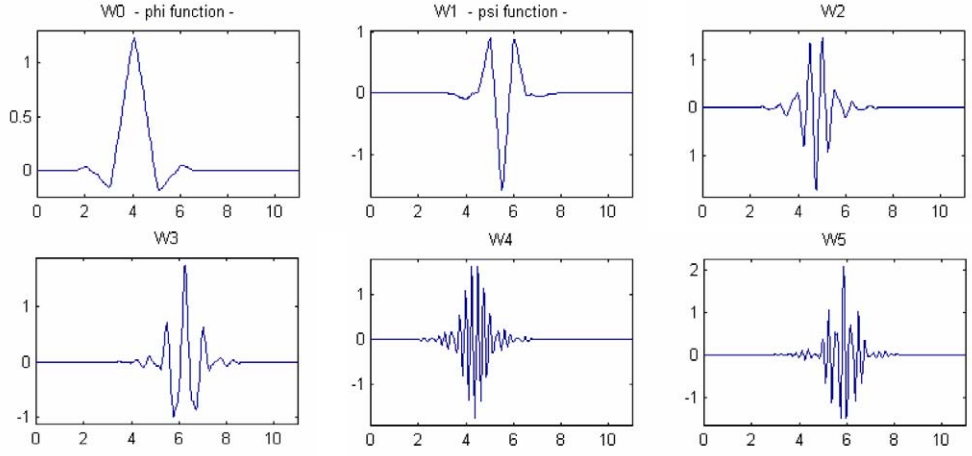


Fig. 13. Example Coiflet wavelet packets (C2) including the generator/scaling functions (W0/1) and higher order members.

tomised biorthogonal wavelets based on the 'lifting scheme' (Sweldens, 1996), which could be tuned to the morphology of specific features or feature characteristics. Such methods could be used go produce customised wavelets for detecting specific features such as calcifications, which have quite well defined image statistical characteristics.

6.4. Extension of wavelet packets to images (2D)

Our treatment of wavelets, wavelet packets and scale-space construction via information cost functions (best basis analysis) has so far been limited to the basic 1D case. However, for analysing images, these concepts need to be extended to higher dimensions. Fortunately, most conjugate quadrature filtering algorithms extend quite neatly to higher dimensions via the tensor product of 1D basis elements:

$$b(x) = b(x_1, \ldots, x_d) = b_1(x_1) \cdot \cdots b_d(x_d). \tag{24}$$

This tensor product basis approach is termed 'separable' because they can be factored across integrals to obtain a sequence of d one-dimensional problems by treating each dimension individually. The equivalent low- and high-pass filters H and G in the definition of the wavelet packet decomposition are quadrature filters with corresponding dual (adjoint) filters that can be used for reconstruction in the same as the 1D case. Therefore the tensor product approach to extend the wavelet packet analysis to higher dimensions naturally involves the tensor product of 1D sequences (filter coefficients), h and g. The tensor product of quadrature filters on the bivariate sequence $u = u(x,y)$ is defined as

$$\begin{aligned}(H \otimes G)u(x,y) &= \sum_{i,j} h(i)g(j)u(2x-i, 2y-j) \\ &= \sum_{i,j} h(2x-i)g(2y-j)u(i,j).\end{aligned} \tag{25}$$

Similarly, for the adjoint (reconstruction) case, the tensor product $(H \otimes G)^* \equiv H^* \otimes G^*$ is defined by

$$(H^* \otimes G^*)u(x,y) = \sum_i \sum_j \bar{h}(2x-i)\bar{g}(2y-j)u(i,j). \tag{26}$$

These definitions generalise to an arbitrary dimension d for a filter operator $F^i \in \{H,G\}$ for $i = 1, \ldots d$:

$$\begin{aligned}&(F^1 \otimes \cdots \otimes F^d)u(x) \\ &\overset{\text{def}}{=} \sum_{k_1} \cdots \sum_{k_d} f_1(k_1) \ldots f_d(k_d)u(2x_1 - k_1, \ldots, 2x_d - k_d \\ &= F^1_{x_1} F^2_{x_2} \ldots F^d_{x_d} u(x).\end{aligned} \tag{27}$$

A visualisation of Eq. (27) for $d = 2$ is shown in Fig. 14.

This approach is analogous to subband coding, a relatively conventional approach to filter decomposition. In terms of implementation, this means that the image is convolved (and decimated) with the quadrature filters in each dimension in a 'quadtree' fashion. As an aside note, the actual implementation of tensor wavelets requires some care with respect to the issue of mixing scales via tensor products (Mallat, 1998).

6.5. Feature extraction

Now that we have defined our notion of a scale-space in terms of 'best-basis' ordered reconstructions from wavelet packets, it is possible to perform feature extraction. The extraction process involves region growing the scale-space and identifying regions that persist through many levels of scale.

Unlike the 'extremum' stack concept of Koenderink (1984) or the conventional CWT scaleogram (Holschneider, 1995), our scale-space construction does not guarantee continuity through the entire stack by virtue of our use of an information cost function to order the wavelet packet construction. Therefore we are interested in regions that may be well localised across several scales, but not necessarily in a congruent fashion across *all* scales. The most significant n regions are treated as 'features' and are labelled by the highest level of scale reached by region growing/merging. A visualisation of this effect is illustrated in Fig. 15 where a small portion of an X-ray image containing microcalcifications (sharp, high-intensity features) is visible as distinct regions propagating through the stack.

The region growing starts with a collection of seed pixels at the low-scale part of the stack. A merging operator is applied to adjacent pixels l_{i+1}, l_{i-1}, l_{j+1}, l_{j-1} and l_{k-1}, l_{k+1} in the scale-stack volume $V(x,y,\lambda)$ using

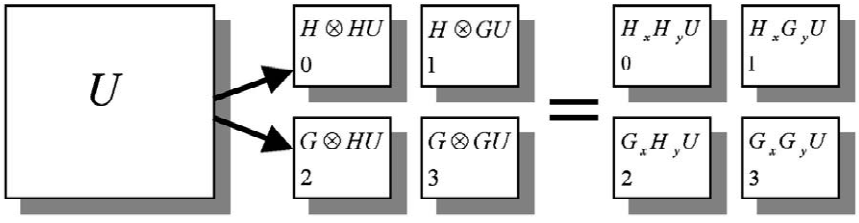


Fig. 14. The implementation of wavelet packet analysis for $d = 2$ is performed by directional convolution and decimation for each dimension of the image.

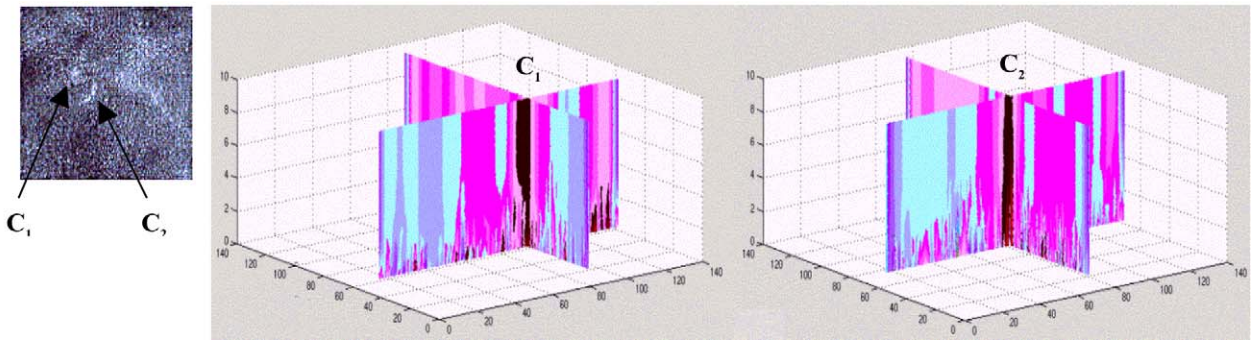


Fig. 15. A small section (140×140 pixels) of a 30-μm digitised mammogram containing two calcifications is decomposed into the wavelet scale-space. The calcifications are detectable as a well-localised propagation through the sliced scale-space (C_1 and C_2). Note: the scale-space has been inverted to better show the shape preservation from coarse to fine scales.

$$\Delta E = \frac{L_i L_j L_k}{L_i + L_j + L_k} \left\| \frac{\sum_{n\in X} V(n)}{L_i} - \frac{\sum_{n\in Y} V(n)}{L_j} - \frac{\sum_{n\in \Lambda} V(n)}{L_k} \right\|^2 - \Gamma \tag{28}$$

and

$$\Delta E \begin{cases} < 0, l_i, l_j, l_k & \text{merge,} \\ \geqslant 0, l_i, l_j, l_k & \text{unchanged,} \end{cases} \tag{29}$$

where

$$L_i = \sum_{n\in l_i} L(n), \quad L_j = \sum_{n\in l_j} L(n), \quad L_k = \sum_{n\in l_k} L(n). \tag{30}$$

$L(n) = 1$ and L_i gives the vector of intensities of l_i. This process is iteratively repeated for a fixed threshold value Γ until no more merges are possible (i.e., the region is completely grown). Γ is then increased until either the desired number of regions is produced (the m most significant regions) or the number of regions converges. In our experimentation, we have found that it is better to use fewer good quality landmarks, therefore this number m can either be controlled interactively by the user (typically $m < 10$) or the landmark selection criteria (described in the next section) used to reject features that have less significant scale localisation. This results in a segmentation of the scale-space that can be visualised as an overlay on the image under analysis (Fig. 16).

Once the regions have been segmented they are fitted with a cubic spline contour and the image labelled into generalised areas of 'region of interest' and 'noise'. The enables a histogram of pixel values to be assigned to 'features' or regions of interest (ROIs; up to a certain scale) and 'noise', which has a different statistical distribution of intensities. With these estimations for feature and noise, we utilise the region competition algorithm of Zhu and Yuille (1996) to slightly refine the contour boundaries of the detected features, as a small amount of smoothing has transpired during the approximation process. This refinement step is not strictly necessary, al-

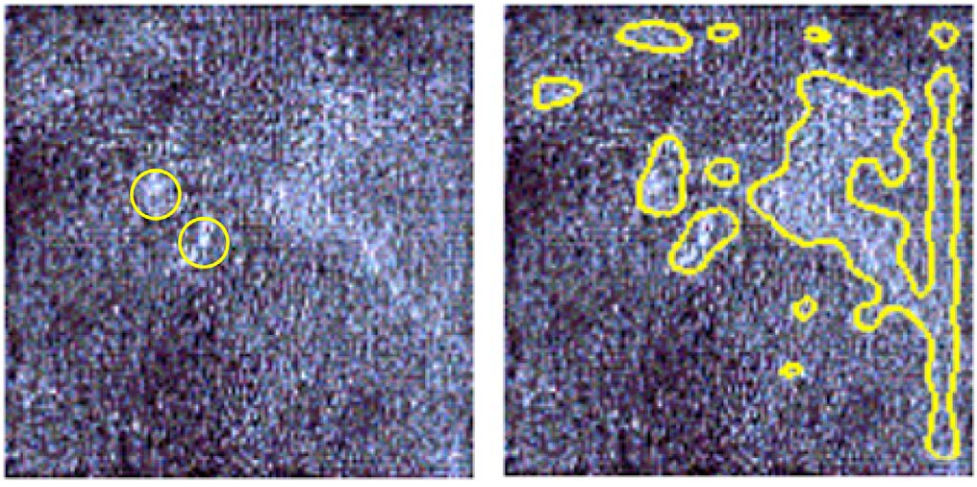


Fig. 16. Segmentation results for the small region of a mammogram illustrated in the previous figure. There are several features in this image (left) including a small curvilinear structure and two clusters of calcifications (circled). The segmentation (right) has produced several regions of interest whilst being relatively impervious to the noisy characteristics of the image.

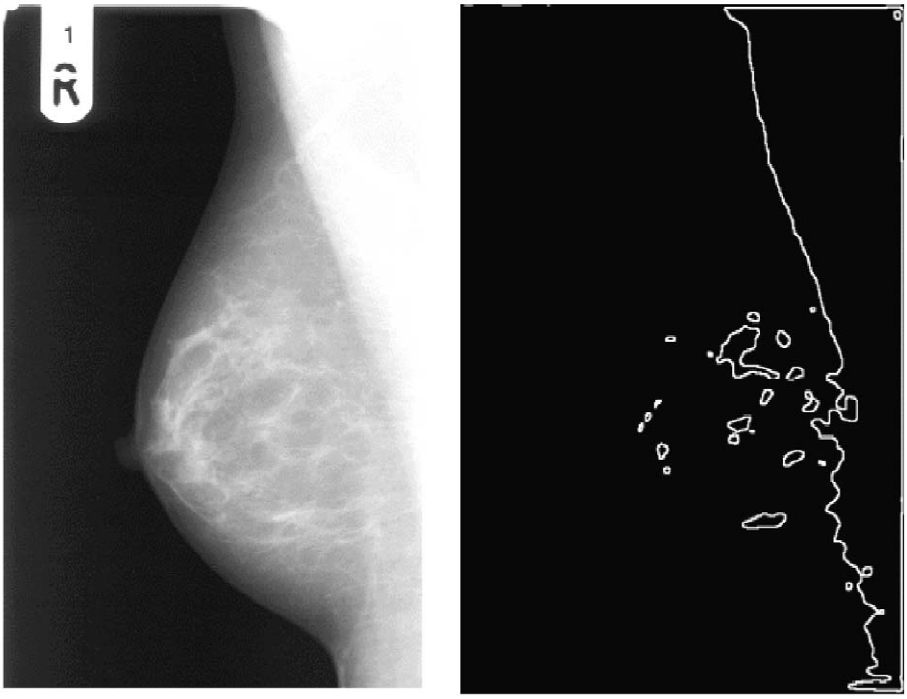


Fig. 17. An example of wavelet-based feature detection and segmentation of a mammogram of a patient with very dense breasts. This is a good example as it shows image scale-space can be used to produce a limited number of high quality landmarks representing visually salient features (breast edge segmented earlier in the process).

though it produces feature contours that have a higher degree of complexity and better reflect the actual morphology (for example the fine spiculations of a tumour). An example of a completely segmented mammogram image is shown in Fig. 17 (minus the breast edge which is segmented earlier). It is this collection of features that is used as the basis of the internal landmark matching.

In terms of actually defining landmarks for the RBF deformation, we ignore the pectoral muscle region and the radiological tag and since our features are based on closed contours grown from the scale-space stack, RBF landmarks are based on the centroid of each feature. Clearly this introduces a degree of error into the registration; however, the approximation framework largely compensates for this in the completion registration step and will be also influenced by the number of landmarks used.

7. Landmark matching

Given a set of segmented landmarks (represented in the deformation framework by a feature centroid), we compute a correspondence measure between features in the registration pair on the basis of the following criteria:

- *Scale localisation* (also used to parameterise Eq. (31)). This reflects the comparative morphology of landmarks as features with similar textural characteristics will be closely related in wavelet scale-space.
- *Orientation.* We compute a simple orientation measure on the feature contour to adjust the sensitivity of the match to ‘blobs’ and curvilinear structures.
- *Feature area.* This does not have the same meaning as the scale localisation parameter, although there may be some overlap in the case of large features. The objective is to ensure that spatially localised features are not matched to larger features. This is parameterised by a fixed threshold (we have found that allowing up to 30% variability provides a reasonable result).
- *Relative motion.* Motion between the landmarks in the undeformed and partially registered (via boundary curvature) data. This is illustrated in Fig. 18, where the relative motion of landmarks is included as a weighting factor in the match.
- *Neighbourhood localisation.* A window proportional to the amount of displacement of the transformed internal landmarks from the boundary registration is used to limit matches to a ‘window’ or neighbourhood.

All of these parameters are used to drive a simple spatial searching algorithm that is conceptually similar to the iterated closest point method (ICP) (Besl and McKay, 1992; Declerck et al., 1996). Essentially these criteria are used as the basis of a match rejection filter to ensure that landmarks classified as a ‘match’ have similar spatial properties. The search characteristics are then also adjusted on the basis of how landmarks have demonstrated a change in geometrical correspondence as a result of the initial boundary deformation—in other words, how the initial registration process improves alignment between internal

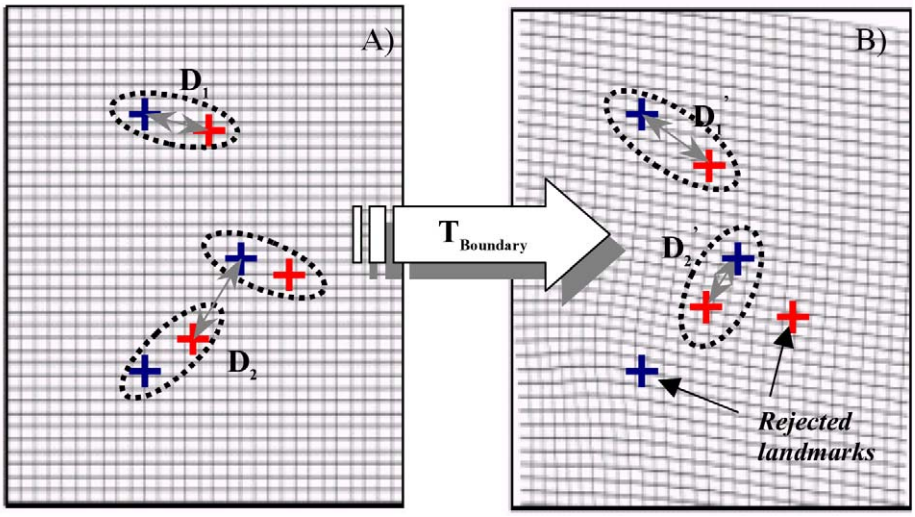


Fig. 18. This figure illustrates how the boundary registration is included as rejection criteria for the matching algorithm. (A) The pre-registration relationship between three points in pair of images to be registered (transformation = identity) with three matches classified on the basis of a closest-point search. D_1 and D_2 represent distance measures between candidate pairs of points. After the boundary registration (B), the candidate pair corresponding to D_1' is maintained, despite a small change in the distance measure (thresholded). However the remaining landmarks are re-matched on the basis of the transformation T_{Boundary}.

features. This partial registration search refinement algorithm is outlined in Fig. 18.

8. Completion of registration using internal landmarks

8.1. Overview

The last step in the registration process is to include both the boundary curvature-based landmarks and detected internal landmarks in the registration process (as outlined in Fig. 1). However, in this stage of the registration process, we use an approximation rather than an interpolation approach to the elastic deformation. This is to account for some of the inaccuracy of centroid-based landmark delineation as well as to produce a smooth deformation that honours landmarks to which we ascribe different degrees of confidence by virtue of wavelet scale localisation properties.

8.2. Approximating radial basis functions

For the boundary registration case, we have relatively high confidence in our landmark selection since it is based on a robust geometrical measure (i.e., curvature) that has relatively little ambiguity. However, in the case of additional internal landmarks, it is important to consider the level of confidence we can place on our selection and evaluation of corresponding features.

In the wavelet feature detection and subsequent feature matching, we have chosen a set of criteria that limits the error in search for feature correspondence. In particular, we restrict our matching on the basis of features with similar scale localisation in the wavelet 'feature stack'. This effectively means that we are not interested in matching features that change size (due to differences in breast compression, for example) or in the case of multimodal fusion, features that may have a different spatial representation between modalities.

To implement this in a registration algorithm in such a way that landmark localisation errors can be accommodated, it is necessary to weaken the interpolation condition in such a way that the resulting transformation is a compromise between smoothness and data adaptation. To achieve this, we define a new functional (refer to Eq. (4)) $J_\lambda(t)$ to be minimised as a function of λ. This λ is a regularisation parameter that controls the balance between smoothness and approximation (Rohr et al., 1996):

$$J_\lambda(f) = \sum_{i=1}^{n} \left| \frac{q_i - f(p_i)}{\sigma_i^2} \right|^2 + \lambda \cdot J_2^2(f). \tag{31}$$

The scale localisation information is used to control the σ_i^2 terms that have the effect of weighting the smoothness in the deformation by the saliency of the landmark. In this way, greater deformation influence can be attributed to features that have stronger spatial localisation. The solution process is very similar to the interpolation problem since only the K submatrix in Eq. (9) needs to be changed:

$$K' = K + \lambda \cdot V^{-1}, \tag{32}$$

where $V = \text{diag}\{1/\sigma_1^2, \ldots, 1/\sigma_n^2\}$. A simple example of the effect of varying σ is shown in Fig. 19(C), where the relative weighting of landmarks produces a smooth approximation.

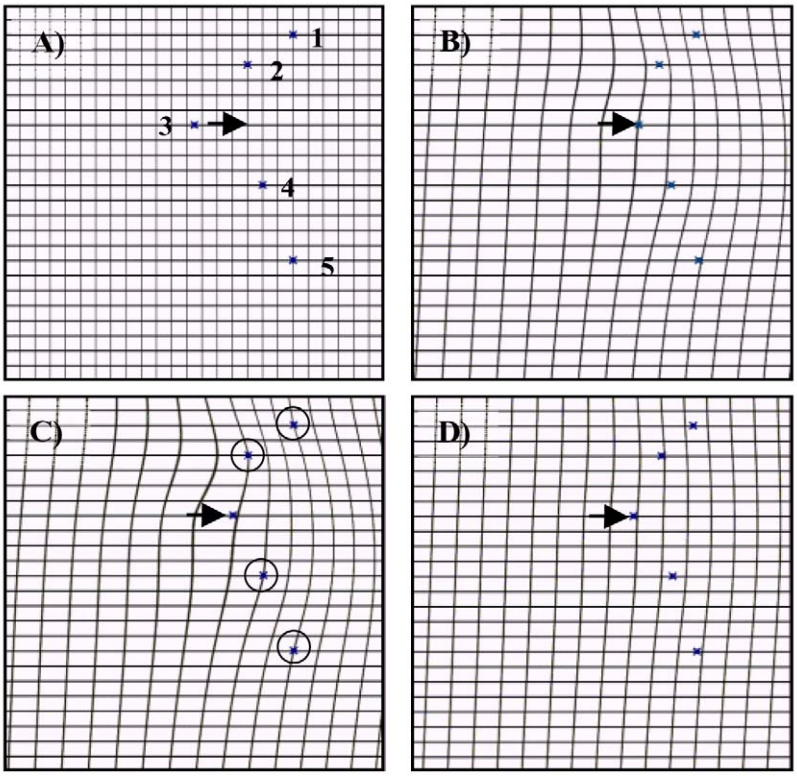


Fig. 19. The regularisation approach of the second step of the registration process. (A) Contains five landmarks (1–5) in the undeformed image. In (B), landmark ≠3 is moved and the resulting deformation approximation calculated assuming that the σ weighting for the remaining landmarks are five times greater in the transformation calculation. (C) The case where landmarks ≠1,2,4,5 are fixed and landmark ≠3 moves. The transformation interpolates all landmarks equally, with a highly deformed result (i.e., the elastic deformation is too global despite the fact that only one landmark moved). (D) The case where all landmarks are used in the approximation process, but the λ set (high) so that the transformation is smooth.

9. Clinical applications

9.1. Overview

Before presenting the results of data fusion for a number of patient studies, we present three particular cases that motivated this work. These case studies are quite specific in the sense that they relate to particular pathological conditions; however, the utilization of our data fusion framework can be generalised to difficulties with imaging the breast in certain circumstances. These studies are also of interest from a validation perspective, because in each case, a wire localisation or needle biopsy was used to confirm the location of the pathology.

9.2. Case ≠1: mammographic imaging of women with dense breasts

There is little contention about the screening value of mammography in the early detection and diagnosis of breast cancer. Mammographic screening has greatly improved the mortality statistics (Fletcher et al., 1993; Nystrom et al., 1993; Verbeek et al., 1984), while emerging technologies such as computer-aided detection (CAD) combined with full-field digital mammography have the potential to further improve the effectiveness of screening (Karssemeijer et al., 1998; Highnam and Brady, 1999).

However, for younger women (typically equated with pre-menopausal), women with very dense breasts and patients who have radiologically exhibited tissue regeneration as a result of hormone replacement therapy, mammography has been proven to be less reliable (Kavanagh et al., 2000; Leung et al., 1997; Litherland et al., 1997). There is even some evidence to suggest that MRI-based screening may be a more appropriate first-line approach to accurate detection and diagnosis for these cases (particularly if a patient falls into a high-risk category) (Warren and Hayes, 2000).

In the first example, a 47-year-old woman presented with a lump that was locally painful and marginally palpable under a clinical examination, but with diagnostically ambiguous mammograms due to the density of the breast tissue. Two attempts at core biopsy, both with mammography and with ultrasound guidance, were unsuccessful in localising the lesion. As a frustrated attempt to obtain more information, the patient was referred for a CE-MRI scan. The results of the mammographic imaging and a relevant slice of the corresponding MRI acquisition are shown in Fig. 20. In this case an extended CC mammogram was taken in the hope of spotting a lesion

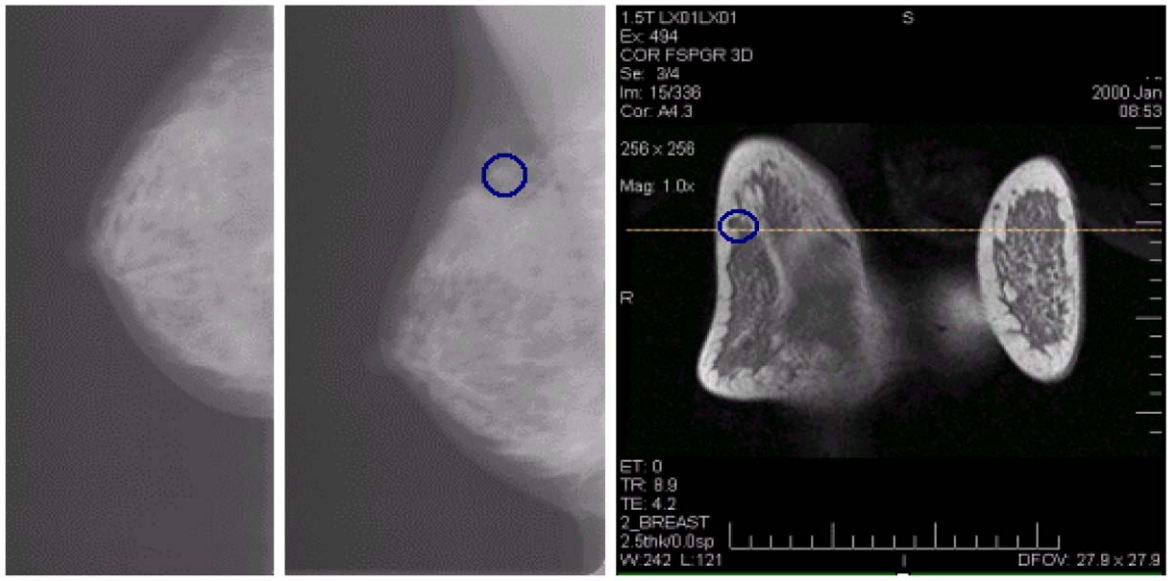


Fig. 20. Initial assessment of the mammograms with the suspected ROI highlighted in the lateral mammogram (centre) along with a corresponding region in a slice of the CE-MRI volume (right). No lesion was obviously apparent in the CC mammogram (left).

closer to the axilla as the conventional CC did not reveal any suspicious regions. The clinician's estimation of the ROI in the lateral view (used for biopsy) and the corresponding assessment of the MRI are highlighted.

The lesion is clearly visible in the MRI acquisition as a moderate-enhancement circular morphology. After registration with the corresponding CC (extended) and MLO mammograms (Fig. 21), the location (centroid) of the ROI in the MRI volume was transformed back to the two mammogram views and presented to the clinician as a marker (Fig. 22).

The projection of the pathology is quite successful in the lateral view, although the CC view (minus the bright region corresponding to the pectoral muscle in the MRI which is not present in a CC mammogram) is obscured by an accumulation of glandular enhancement (a small amount of grey-level rescaling helps somewhat). The inverse transformation of a clinician-selected point in the MRI volume produced a slightly different suggested location of the pathology when compared with the clinical assessment of the original mammograms. In this case, the patient was imaged retrospectively using MRI and a localiser inserted. The patient mammograms were then retaken with the localisation wire to confirm the registered estimation of the ROI. A subsequent biopsy tested positive for malignancy. The actual location of the pathology in the ex-CC and MLO view is shown below.

From a clinical point of view, the patient management was difficult in this case. A long period elapsed between the initial examination and the final positive needle biopsy. Though this was partially due to the number of imaging stages required to confirm the diagnosis, the ambiguity of indicators present in the radiological assessment was a major contributing factor. Naturally, this created both clinical pressure and patient stress that might have been alleviated had the diagnostic indicators been less complex. Although the additional MRI and post-localisation mammograms were possibly an unnecessary process, a robust

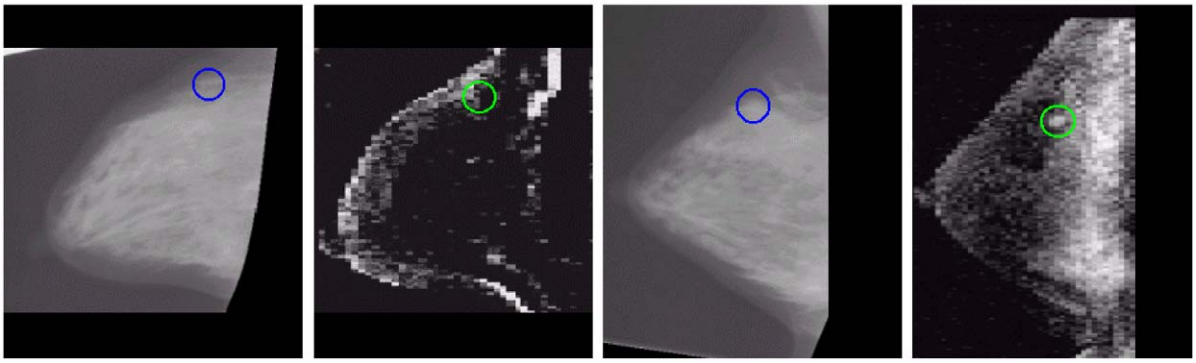


Fig. 21. Registration results between both mammogram views and the MRI projection with corresponding ROI location highlighted.

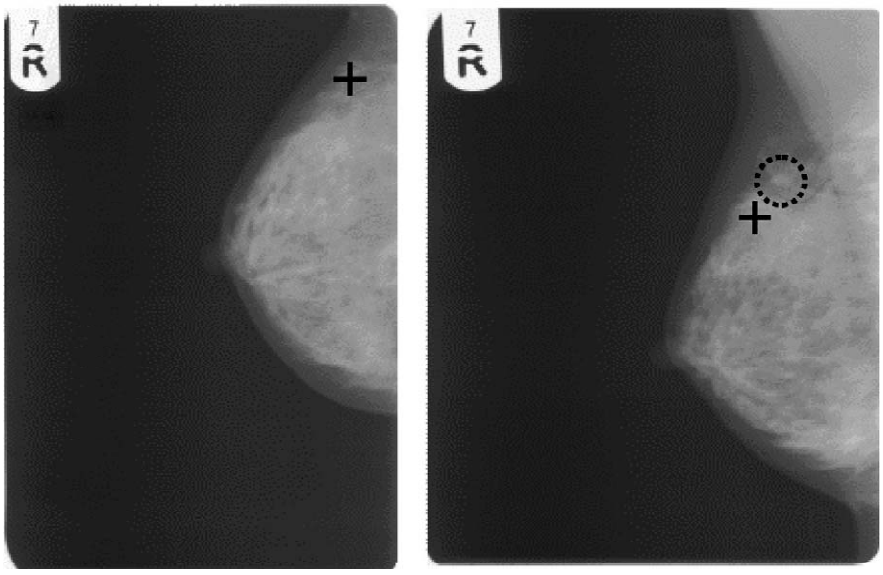


Fig. 22. The initial clinical assessment of the extended CC mammogram (left) did not reveal any suspicious regions of interest. The MLO mammogram (right) was determined to have a slightly denser region (circled) but did not reveal any pathology with needle biopsy. After registration with the CE-MRI projections, the centroid of the region visible in the MRI was projected back onto the mammograms (indicated by +). Biopsy of this region tested positive for malignancy.

implementation of our fusion approach could eventually facilitate this kind of comparison at the time of the MRI acquisition.

9.3. Case 2: ductal carcinoma in situ (DCIS) characterised by microcalcifications

Ductal carcinoma in situ (DCIS) is a commonly occurring form of breast cancer that can be difficult to assess mammographically. A clear indicator of DCIS is often the presence of calcifications with a 'jagged' morphology (Lanyi, 1996; Wolfe, 1982), which are usually quite clearly visible in at least one of the mammogram views. A difficulty with DCIS is the often impalpable nature of the tumour during clinical examination. From a surgical perspective, therefore, 3D localisation and visualisation is very useful. This is because of the high dependency on the palpability of the tumour to guide the resection process. An obvious way of localising the tumour in 3D would be a CE-MRI acquisition.

Unfortunately, the sensitivity of CE-MRI to DCIS is very low due to the relatively small amount of tumour vascularity and hence the contrast-enhancement profile of the pathology is usually too subtle for reliable detection (Rankin, 2000). However, if the locations of the calcifications in the mammogram views could be transformed to the spatial coordinates of the MRI volume, it might be possible to prompt the radiologist with regions of interest that should be more clearly scrutinised for suspicious (although subtle) enhancement and morphology. Of course, an additional complication with imaging this kind of pathology using MRI is the relatively small size of the cancer and hence the limited resolution of MRI (and partial volume effects), which further compounds the issue.

In this example, a young (26 years of age) woman presented for a routine mammogram screening which revealed three distinct clusters of calcifications largely visible in both views. Normally, the age of the patient would preclude a mammogram, however she had already had an occurrence of ductal carcinoma 4 years previously, which was treated by a wide local excision in the upper (outer) quadrant of the breast followed by radiotherapy. Previous surgery also adds complication to mammography imaging as surgical scarring effectively masks pathology (scar tissue has unfavourable X-ray projection characteristics) and therefore the patient was sent for a follow-up MRI. However, the clinicians were interested to know the correspondence between the calcification clusters visible in the mammogram (despite the dense tissue) and the CE-MRI acquisition. Fig. 23 depicts 4 years of mammograms (both views) that have been registered using the technique described in (Marias et al., 2000). In this way, it is possible to compare the changes in tissue density due to both the surgical treatment (the excision was followed by a breast reconstruction using muscular tissue which is clearly evident in the sequence) and possible evolution of pathology.

This first stage analysis of the mammographic history of the patient was a useful exercise. Firstly, the surgeon felt that the calcification cluster located near the previous excision site (cluster labelled '1') was a relatively low-risk of cancer since surgical scarring may produce a slight dusting of calcifications. Certainly the shape and distribution characteristics of the microcalcifications were

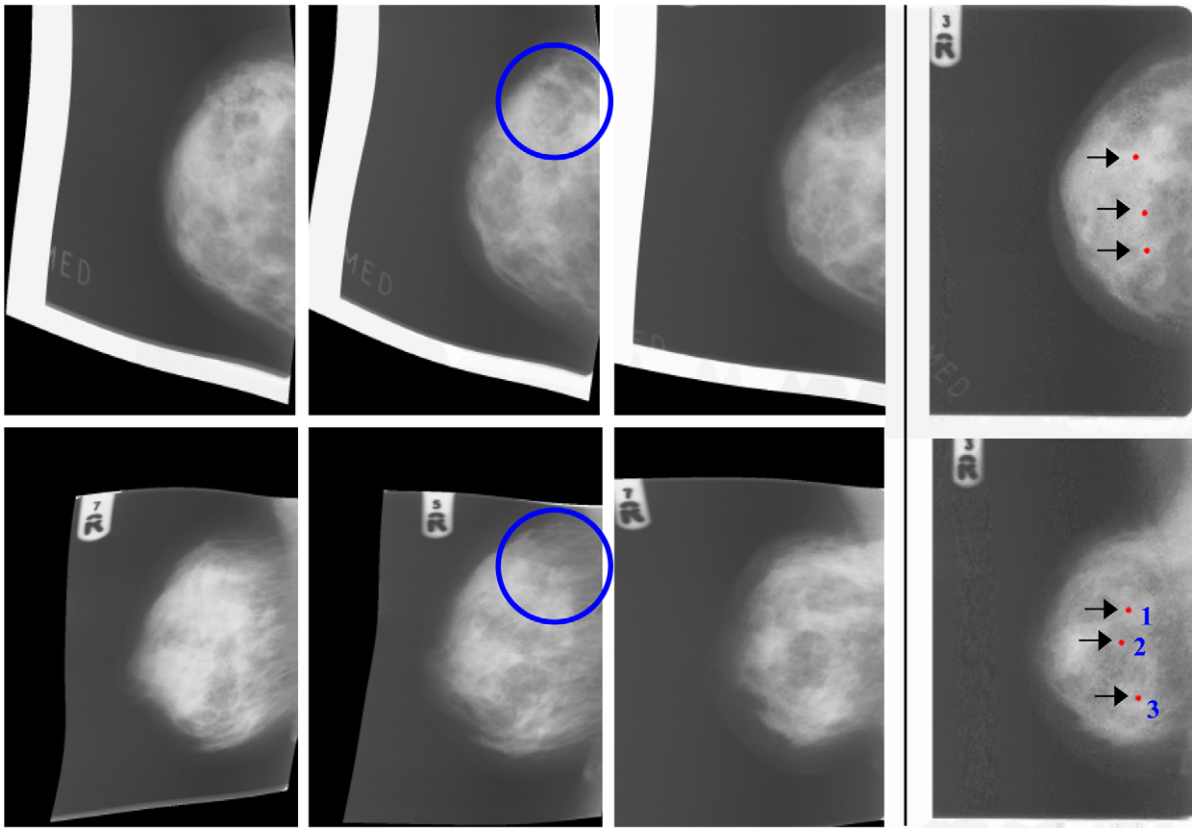


Fig. 23. Four years of patient mammograms registered using the techniques described in (Marias et al., 2000). At year 4, three clusters of calcifications are visible in both the CC (top row) and MLO (bottom row) mammograms. Comparing the location of cluster '1' to the surgical excision area (circled regions in the mammograms of year 2) suggests calcification due to surgical scarring rather than malignancy.

consistent with this expectation. The second cluster of calcifications (labelled as '2') was somewhat contentious due to the location, particularly with respect to the nipple plane of the breast. The remaining cluster was morphologically the most pathologically indicative of DCIS.

The two mammogram views were registered to the CE-MRI volume in the manner described previously and the location of the calcifications were colour coded and presented to the radiologist in a 3D slice-based visualisation of the MRI volume. This is shown in Fig. 24 where the 'trouble' case of calcifications ('3') is visualised as a red spherical region of interest (indicating the approximate extent of the cluster).

By providing a set of visual prompts to the radiologist we had hoped to highlight regions that required additional scrutiny. Since one expects the contrast-enhancement characteristics of the calcifications to be very subtle in the CE-MRI acquisition this provides a useful additional piece of information to the radiologist. Upon re-examination of the MRI contrast sequence, the radiologist concluded that there was a region of slightly higher enhancement at the indicated position of the third cluster, as indicated in Fig. 25.

Although the enhancement characteristics shown in the previous figure are subtle, there was enough quantitative evidence (from the contrast profile) to indicate that the suggested location of the calcification cluster in the reference frame of the MRI was potentially significant. Biopsy of the calcification cluster proved positive for DCIS.

An area of significant debate in breast MRI is the extent to which DCIS enhances with gadolinium DTPA. Some radiologists suggest that there may be a subtle enhancement characteristic that can be used to detect DCIS. We are interested to pursue this as a larger study to see whether a definitive quantification of DCIS enhancement can be obtained. This could lead to the development of software tools that not only aid the localisation of mammogram-visible microcalcifications in the MRI acquisition, but also provide analysis tools that help the radiologist localise DCIS-characteristic enhancement. In the interim, we use calcification localisation in conjunction with some simple segmentation techniques based on region-growing local voxel contrast characteristics to produce visualizations of DCIS that we feel are useful to both the radiologist and ultimately, the surgeon. An example of such a visualisation (based on a sagittal reslicing of the breast of interest) is shown in Fig. 26.

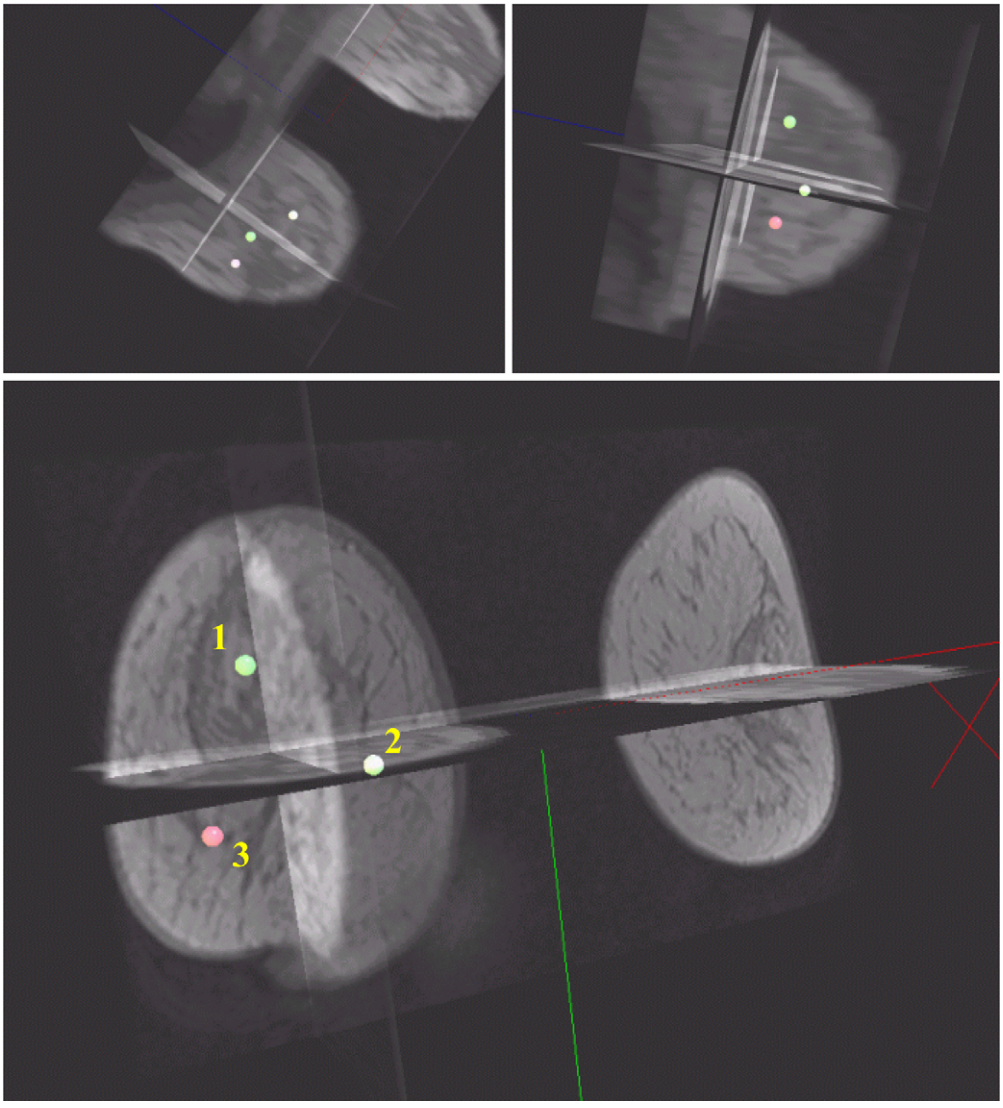


Fig. 24. 3D visualisations of the location of the three calcification clusters in the reference frame of the MRI volume (labelled according to the previous figure). This figure (and several other figures in this paper) is available in colour, see the on-line version.

9.4. Case ≠3: multi-focal disease and correlated pathological indicators

One of the difficulties with mammography-based assessment is the difficulty of determining the extent of multi-focal disease. Sometimes a lesion visible in a mammogram is only one of several present or in severe cases, bilateral lesions are not evident. With respect to MRI, one of the major trade-offs is spatial resolution versus temporal sampling of contrast uptake. Radiologist colleagues assume that there is always a possibility of multi-focal disease, particularly in patients who have exhibited recurrences. Therefore despite the fact that improved temporal resolution would give better (more discriminating) contrast enhancement information, the clinical value of including both breasts in the examination is simply too important.

In terms of improving the bi-directional flow of information between a mammographic and an MRI-based examination, comparing multi-focality is a key issue. Multi-focal disease is more likely to be evident in MRI and therefore if we wish to prospectively evaluate a mammogram, data fusion is useful. An example of this scenario is shown in Fig. 27, where a fairly evident mass highlighted by calcifications is present in the mammogram and more than one region of focal enhancement is visible in the CE-MRI acquisition. This can be clearly seen by comparing the pre- and post-contrast MRI images.

In the case of multi-focal disease, there is some benefit in being able to compare regions of enhancement in the MRI volume, with corresponding regions in the mammogram. In this way, fine mammographic detail such as calcifications and spiculations can be compared with focal

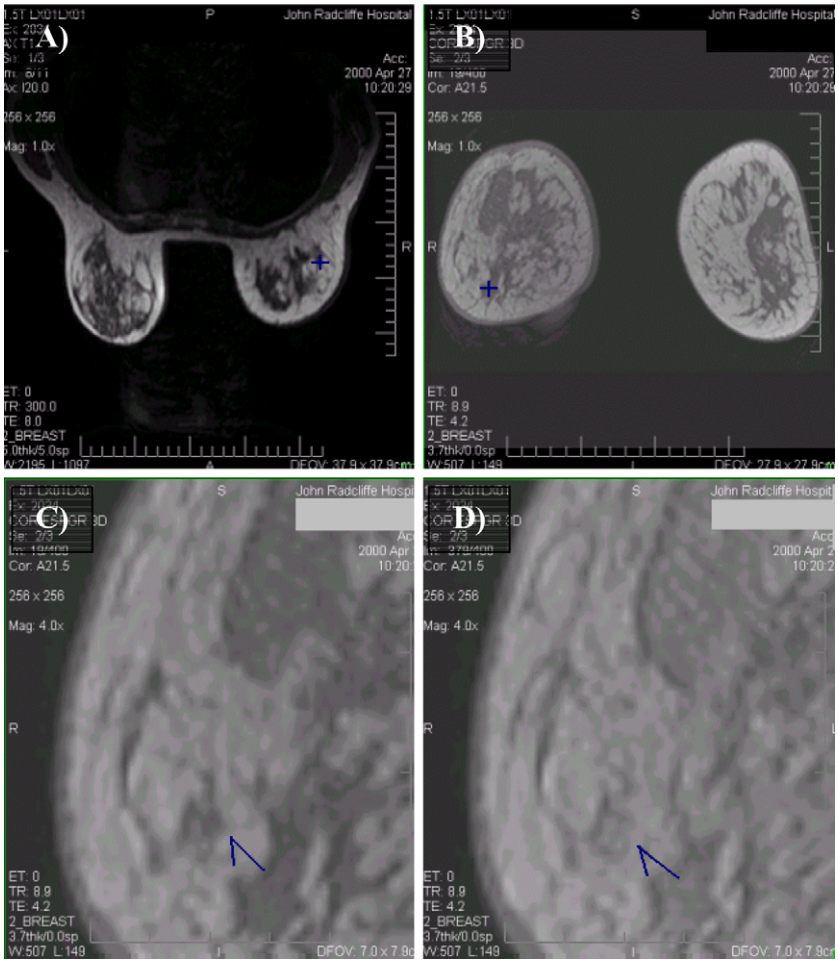


Fig. 25. (A,B) The location of calcification cluster '3' in the corresponding coronal and transverse slices. Magnification of this area in the coronal slice for (C) pre- and (D) post-contrast indicates a slightly enhancing region.

enhancement in the MRI volume, enabling the radiologist to evaluate pathology both in terms of the structural characteristics for which mammography is very detailed, combined with the functional representation of disease in CE-MRI. For the case shown in Fig. 28, we illustrate a selection of points in the MRI volume that were selected by the clinician to compare the enhancement of suspicious regions with normal parenchyma contrast uptake. The displayed points are actually selections across several slices displayed as a cumulative overlay to provide the clinician with a sense of the spatial distribution of points within the reference plane of the current slice. The corresponding points have been registered to the mammogram views and are displayed in the lateral mammogram for comparison. The blue crosses indicate the two regions of enhancement seen in the MRI volume.

This is an interesting result in the sense that it illustrates how difficult it is for the radiologist to mentally perform this comparison. The combination of landmarks selected across several MRI slices and the relative complexity of the internal deformation of the mammograms during non-rigid registration illustrates why radiologists appear to generally be unable to perform this comparison manually. In fact, the radiologists we work with do not usually refer to the mammogram when assessing multi-focal conditions as it is considered possibly misleading. We believe that the registration framework presented in this article is a step towards making this comparison possible.

As a final note to this usage scenario, we emphasise that CE-MRI provides a functional representation of the structural detail in a mammogram. Mammographic assessment of large, poorly differentiated carcinomas, such produces relatively little detailed information about the necrosis of the tumour. Although it would be useful to be able to learn something about the detailed pathological structure of the tumour, this is often not possible with a mammogram. On the other hand, the MRI contrast-enhancement characteristics provide valuable information about the vascularity and necrosis of a tumour. The projected contrast enhancement from the MRI acquisition can produce a representation of the pathology that can be compared against detail visible in the mammogram. As shown in Fig. 29, this can be used to

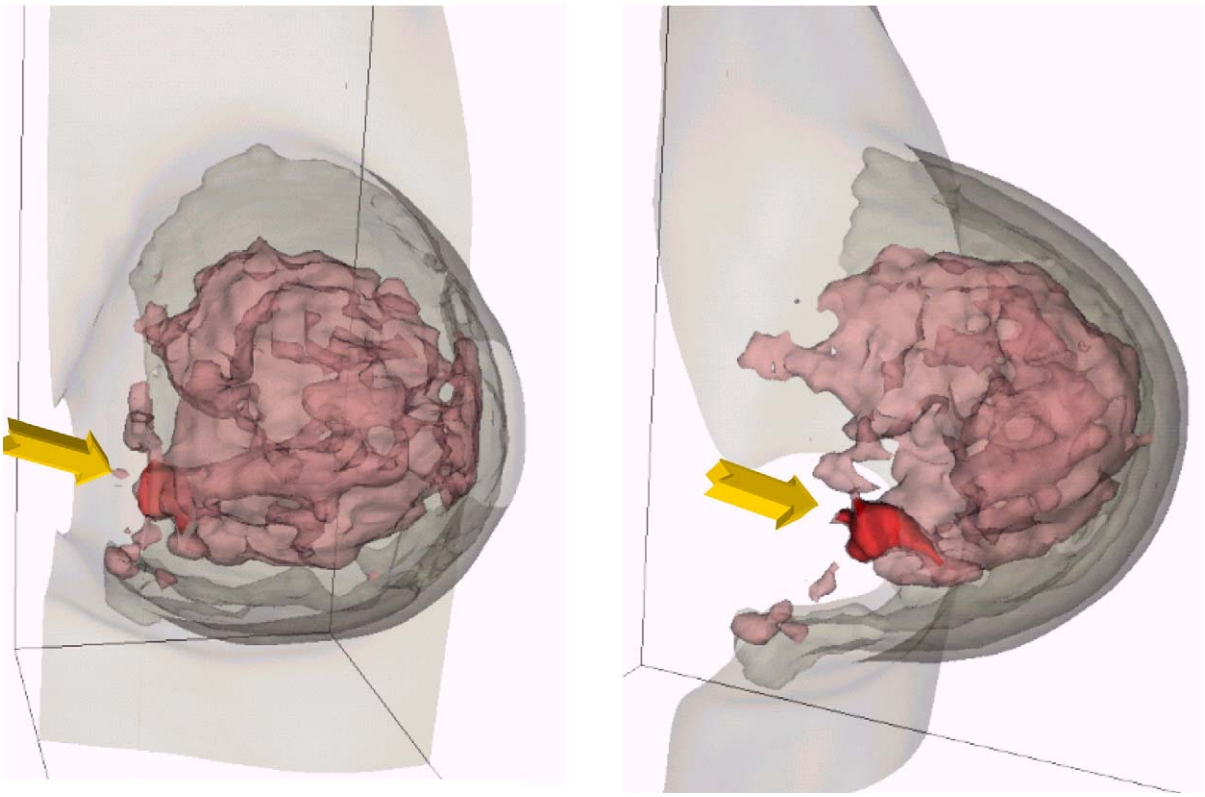


Fig. 26. A rendering of the MRI volume (post-contrast sagittal acquisition) with the region corresponding to calcifications in the mammograms indicated by arrows. In terms of understanding the pathology, 3D morphology is much more meaningful than a slice-based representation. Additionally, models such as this may provide the surgeon with a useful location reference during excision.

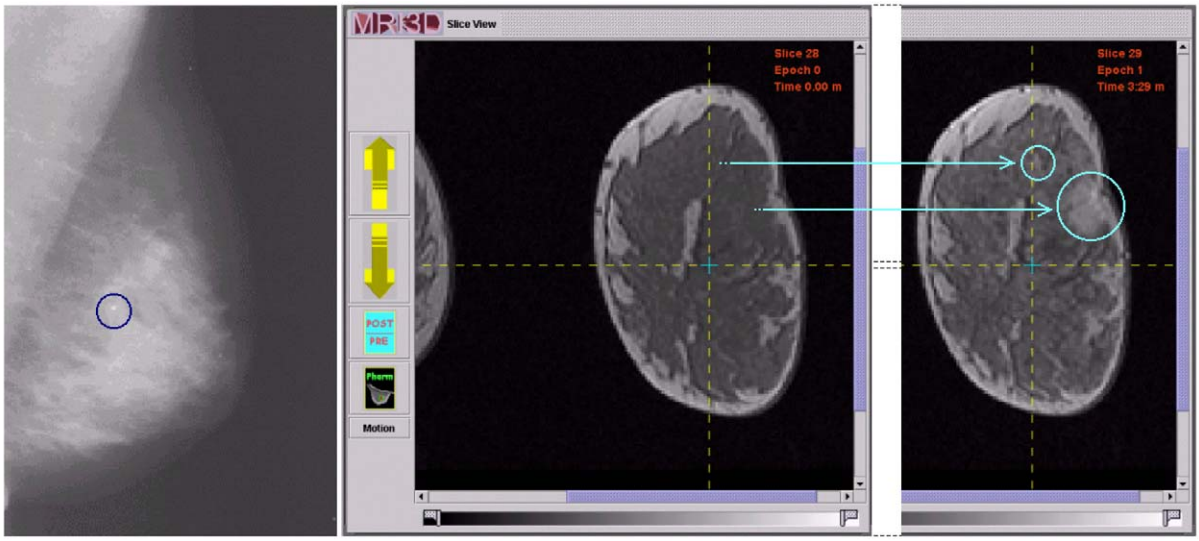


Fig. 27. A suspicious region (indicated by calcification) in the lateral mammogram (left) compared with a pre- and post-contrast slice in the MRI. The enhancement characteristics of the MR images suggest more than one enhancing region (circled).

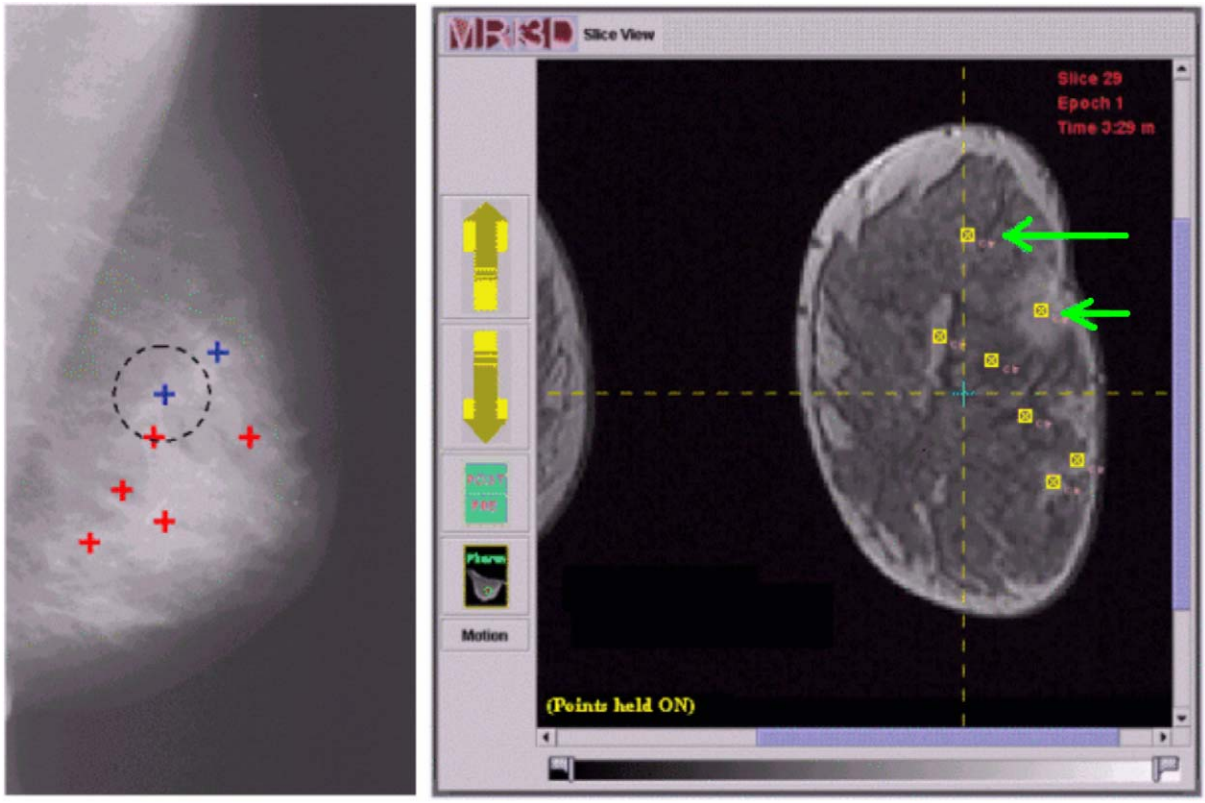


Fig. 28. A collection of points in the lateral mammogram (A) corresponding to points selected across several slices in the MRI volume (displayed in B) as an overlay). The circled region in (A) encompasses the region in Fig. 27, although it is clear from the MRI data that there is more than one enhancing region (blue crosses).

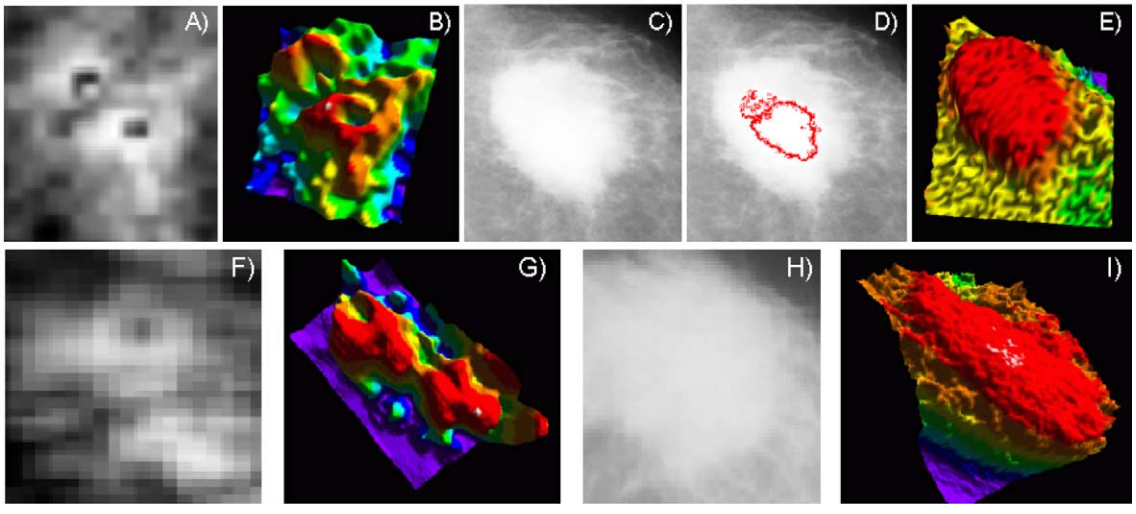


Fig. 29. (A,F) The contrast enhancement projections corresponding to (C,H) in the mammogram (poorly differentiated lesions). In both cases, the MRI projections illustrate small ‘cores’ of lower enhancement that typically correspond to necrosis. Surface rendering of these regions (B,G,E,I) provides an interesting comparison of the shape of the two representations. (D) Simple thresholding and edge-detection of (C) provides a measure of substantiation that there is more than one component to the mass.

compare the spatial distribution of the necrosis of the tumour against the extent of the mass in the mammogram (shown for two patient examples).

10. Registration results

In the previous section, we illustrated through a series of patient cases, the clinical value of mammography to CE-MRI data fusion. We have had the privilege of demonstrating our software tools on a range of pro- and retrospective cases that we believe have made a contribution to the patient management in several instances. As a matter of course, the results of our registration and analysis can only be accepted as a clinical possibility due to the difficulty in validating non-rigid registration algorithms—particularly between different modalities. In a few cases, we have had the benefit of a follow-up study involving image-guided localisation with either mammography or MRI, although this would not be the usual case due to the patient inconvenience and cost of such assessments.

To date, our technique has been applied to 14 sets of patient data with good results. In two cases, we were unable to assess the registration due to the fact that features such as calcifications were not present in both mammogram views. This represents a limitation of our method. In the remaining 12 cases (some of which were presented in (Behrenbruch et al., 2000a,b), we have assessed the performance of this projection/registration approach on a wide range of pathologies with very different structural and tissue properties. To a certain extent, the number of studies available for evaluation has not been restricted by data availability, but rather clinical need for such a comparison (and consequentially, the clinical desire to spend time on follow-up analysis). In most cases, women who have undergone an MRI scan have the necessary mammographic information available to perform data fusion. A larger retrospective study could be performed to enable a measure of validation against clinical cases where there is less ambiguity in the radiological assessment.

Our evaluation of the performance of our registration, although limited, is quite simple. In cases where pathology is visible in either/both of the mammogram views and the MRI volume, or a biopsy/localisation has been performed, we compare the centroid of the backwards transformation from the MRI volume to the two mammographic views on the basis of a lesion (or ROI) centroid. This analysis centroid is then compared against the centroid of a manual segmentation by the clinician. On the assumption that different classes of lesion will result in different deformation results (and corresponding errors) we have grouped each case into 7 species of pathology. These are shown in Table 1. It is not terribly surprising that larger pathologies exhibit a more substantial registration error as the use of these regions as a salient feature for registration is more ambiguous by virtue of our method for producing landmarks for deformation. Similarly, it is sometimes problematic to evaluate cases with calcifications as there may not be corresponding enhancement in the MRI volume for all clusters. Again, these considerations have limited the size of this study.

The set of results illustrated in Table 1 is somewhat difficult to compile as the deformations are calculated in mm, but the voxel dimensions are not uniform across the various MRI acquisitions. Therefore, we have quoted the error relative to a voxel as it better illustrates the fact that in many of the cases, the MRI scans were acquired using transverse slicing and hence the voxel resolution in the MLO projection is poorer. This error trend is quite visible when the MLO voxel error is compared against the CC error. To further understand the relative significance of this error, we compare the estimated error in mm (converted on the basis of the average voxel dimension in both projections) to the average calculated displacement field computed from the registration transformation function (conceptually shown as a circular radius of deformation in Fig. 30).

It can be seen from the chart in Fig. 31, that although the estimated registration error compared with the manual segmentation of the lesion is (in absolute terms) quite large, the displacement necessary to compute the transformation from the compressed shape of the breast in the mammogram to the prone shape of the MRI volume is substantial. On average across all seven classes of pathology, the registration error relative to the size of the deformation is less than 20%. Although this is a preliminary result, it provides an indication of the extent of the

Table 1
A comparison of the registration error in 12 cases and categorised by 7 classes of pathology

Class	Description	CC Error	MLO error	Avg. error	No. of cases
1	Small circular opacity	3.5 voxels	2.5 voxels	3.0 voxels	1
2	Spiculated opacity*	2.0 voxels	5.0 voxels	3.5 voxels	1
3	Infiltrating carcinoma	4.0 voxels	3.5 voxels	3.75 voxels	1
4	Fibroglandular Nodule	1.5 voxels	2.5 voxels	2.0 voxels	1
5	Large, dense circumscribed tumour	3.0 voxels	5.5 voxels	4.25 voxels	2
6	Multiple microcalcifications	2.5 voxels	3 voxels	2.75 voxels	5
7	Fibrous lesion (low enhancement)	2.5 voxels	6.5 voxels	4.5 voxels	1

*Associated with enlarged lymph node.

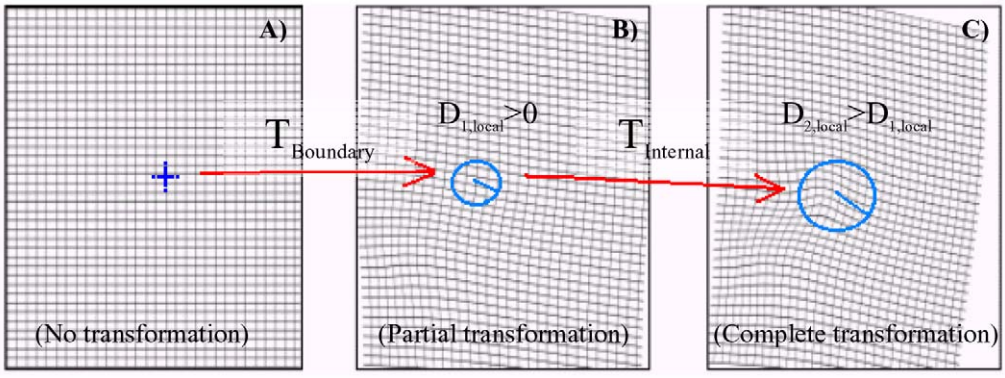


Fig. 30. An illustration of the local displacement field D calculated across the two successive registration steps. The local displacement field is calculated (from the transformation) for a point of interest and is used as a comparison with the registration error for a clinical ROI.

deformation necessary to achieve registration. Once again, however, the extent of the local displacement of a region varies between the two projective views but also depends greatly on the shape and size of the breast and to a certain extent, the slice orientation of the MRI acquisition.

In summary, although not trivial to validate, the maximum registration error we encountered was in the vicinity of 5–6 voxels, which in the case of a transverse MRI acquisition (with 2–3-mm voxels in the MLO projection), this means a maximum error of typically 10 mm (depending on the projection direction). Of course, in cases where the pathology is clearly visible in both modalities, this can itself be a registration landmark—hence in some instances, very good registration results were achieved with a perceived accuracy of 2–3 mm. However, it must be considered that in these cases, the centroid error for larger lesions is also likely to be a somewhat inaccurate reference point as the extent of the visible lesion in the MRI volume will, when projected, underestimate the size of the region in the mammogram.

11. Discussion

11.1. General comments

In this article we have endeavoured to present a novel concept in breast data fusion. To our knowledge, such an approach has not yet been attempted in a clinical application and therefore there are naturally a number of issues that require further assessment. One of the major difficulties is validating a system that has a large number of components and a relatively high degree of complexity. There is little doubt that each component in the process involves a number of assumptions, and whilst we have

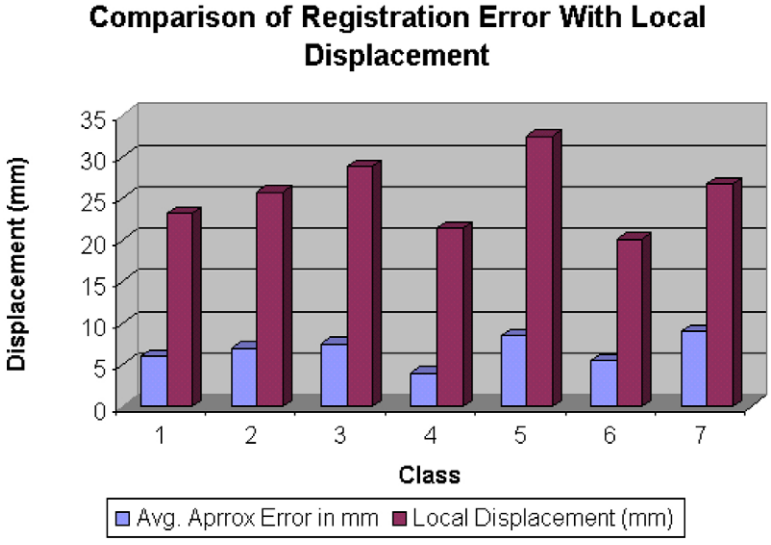


Fig. 31. A graph showing a comparison between the clinician evaluated registration error (with respect to features in the MRI projected back to the two mammogram views) and the local displacement resulting from the elastic deformation. On average, the registration error is less than 20% of the local displacement calculated from the transformation.

striven to minimise the heuristic nature of each step, there is clearly room for improvement.

If such a system is to be clinically implemented and considered useable by a radiologist, the key to success will ultimately lie in the software implementation and interface. The techniques presented in this paper have been presented as automatic; however, in practice there is a need to be able to provide the clinician with interactivity at each step. Not only is this likely to improve the clinician's faith in the registration process, but it enables the 'expert opinion' to influence the quality of the result. To this end, our software implementation of the techniques presented in this paper (as a system) allows user input at each step of the process.

In general terms there are a number of areas that need to be addressed as 'caveats' when implementing this kind of breast data fusion. Some of the obvious issues include:

- *MRI resolution.* Some of the data in our study used slice-based 2D data (i.e. there are slice gaps), which requires volume interpolation prior to projection. The registration error is strongly influenced by the voxel resolution and therefore this is clearly an issue. The good news is that our current MRI protocol involves a true 3D acquisition, which improves the outcome dramatically.
- *MRI slice orientation.* 3D MRI acquisition is usually based around a coronal slicing profile. Although a relatively high number of slices can be acquired (we have examples with as high as 70 slices), this still limits the projection resolution considerably. Ironically, the older transverse acquisitions that provided 256×256 voxels in the CC projection plane produced better results in some cases. In the future, if breast data fusion is required, it may be necessary to take several different slice configurations in order to optimise the resolution of the different projection directions (i.e., a transverse and a sagittal view).
- *Breast compression.* The key idea of this work is that the registration process 'uncompresses' the mammographically imaged breast into the shape of the prone MRI breast. Our model is simple and despite some allowance for internal deformation via the scale-space feature detection, there is room for improvement. Perhaps a combination of the physics-based models (Highnam and Brady, 1999) and a registration approach could be the answer. Alternatively, there may be some value in pursuing a registration process that involves imaging a compressed breast with MRI. There are several interesting possible research directions here.
- *Tissue characterisation.* There are two issues here. Firstly, if the breast is highly involuted in the mammogram views, then the internal deformation characteristics of the registration will be a poor approximation to compression compensation. Fortunately, in these cases the radio-opacity characteristics of the tissues in the mammogram are favourable, and in general are not representative of studies that would require this kind of fusion process. It is still, however, a consideration. Additionally, our MRI projection depends on accurate characterisation of the tissue pharmacokinetics, an issue that certainly warrants additional attention in order to produce 'pseudo mammograms' that provide a meaningful basis for comparison with mammography.
- *Motion correction.* Although we have not discussed this issue extensively in this paper, MRI breast motion correction is a critical issue. In a few cases where we did not correct for motion (as the extent of the pathology was clearly visible under enhancement) the MRI projection was susceptible to cumulative motion artefacts. Whilst this may not detract from a point comparison between views, it significantly affects the robustness of internal landmark detection.
- *Internal landmarks.* Our wavelet-based feature detector is a promising start towards compensating for the internal deformation of the breast. We chose a wavelet framework as it provides a large degree of flexibility with respect to utilizing spatial filters that are morphologically sensitive. In our case, the notion of a feature is essentially a well-localised 'bump', however there are many image characteristics such as curvilinear structures and the outline of the ductal disc that could be included in the registration framework to provide a more accurate internal deformation.
- *Deformation regularisation.* The use of radial basis functions for elastic image registration has proven to be a useful, clinically intuitive and computationally effective way of performing the data fusion. By being able to interact with the landmarks for registration and quickly computing the registration result, the clinician is able to efficiently achieve an acceptable fusion. However, we accept that regularization, particularly in the approximation context, is a complex issue and is currently somewhat heuristic. The notion of scale-controlled smoothness has enabled us to partially reject the influence of small features that may introduce a registration error; however, there is still much work to be done. Future efforts might involve the use of deformation models that pay greater heed to the biomechanical properties of the tissues under deformation. Whilst this would not be a trivial exercise, considerable progress has been made with models that demonstrate soft tissue characteristics that are quite realistic (Cotin et al., 1999; Azar et al., 2000; Picinbono et al., 2000; Kyriacou et al., 1999).

11.2. Future validation

Mathematical and clinical validation of non-rigid registration is a problem that is an ongoing burden to the computer vision community. We have provided a measure of validation against deformation characteristics, clinician segmentation and in a few cases we have had the benefit of

a localisation study. However, future research effort could involve further validation exercises, including the following:

11.2.1. Clinical population

The case studies presented in this article have been motivated by the clinical need for fusion. However (at least in the UK) the majority of patients who present for a MRI exam have the necessary mammogram studies for fusion. It is conceivable that quite a large clinical study could be performed on cases where the radiological 'ground' truth is well understood. The problem with this may be the fact that the lack of ambiguity in the clinical assessment is related to the characteristics of the mammograms (such as involution) and therefore in terms of compensation for internal deformation of the breast, these studies may be of limited value.

11.2.2. Histology

The absolute 'ground truth' for validating non-rigid registration is biopsy and wire localisation. We have had the benefit of evaluating our technique with a few localisation studies, however this would need to be performed on a much wider scale. A possible protocol for this would be MRI-guided wire localisation followed by a mammogram(s) and core biopsy.

11.2.3. High-speed MRI

Earlier in our discussion we mentioned the possibility of compressing the breasts during the MRI acquisition to provide a better geometrical correspondence between MRI and mammography. Various kinds of MRI-imaged breast deformation studies have been performed to look at tissue deformation characteristics from a diagnostic standpoint. There is a steady emergence of near real-time MRI sequences, particularly for applications such as cardiac motion tracking. An interesting approach to validation of breast registration would be to image the deformation of the breast and treat the acquisition as a motion-tracking problem. This could be compared against the kind of transformations obtained via deformable models or registration techniques. A few sample frames of such a sequence are shown in Fig. 32, where a variation on a fast, spoiled GRASS acquisition (20 frames/s) captures the motion of a volunteer compressing the breasts in a transverse fashion. This example was not obtained with a breast coil, but with a cardiac surface coil above and below the patient—however, it demonstrates the concept quite well.

12. Conclusion

We have introduced a process of breast image data fusion designed to aid the radiologist in understanding the pathological relationship between mammography and MRI, particularly when there may be ambiguity in the diagnostic value of either modality. Our framework has been developed to address a fairly wide range of imaging conditions with respect to MRI acquisition protocols and patient tissue characteristics. The results of the techniques presented in this article have been assessed with a study of 14 patients with promising performance in 12 of these cases.

Our data fusion approach incorporates a combination of MRI contrast-enhancement projection (to produce a 'pseudo' mammogram), boundary and internal landmark detection and a two-stage registration process using radial basis functions. The internal deformation of the breast (due to mammography compression) is compensated by the correlation of internal landmarks between the MRI projection and the mammogram, using a novel wavelet-based scale-space analysis. By combining boundary and internal landmark information, the registration approach is able to cater for cases where there may or may not be internal landmarks due to the involution characteristics of the breast.

Finally, we presented an interesting set of clinical cases where our data fusion method was required to improve diagnosis or provide additional pathological information.



Fig. 32. Three frames taken from a high-speed MRI acquisition of breast deformation. In this simple example, the patient palpated her breast (hand location illustrated by the arrows) whilst lying prone in the MRI machine between two cardiac surface coils. The internal deformation is quite clear and could be used (in the context of a motion tracking approach) to evaluate deformation fields produced by non-rigid registration algorithms.

These case studies were motivated by a clinical need to rectify the diagnostic disparity between mammography and MRI, most often in cases where the patient exhibited dense breasts that did not image optimally with mammography. We then extended our set of case studies to a brief assessment of 12 (of 14) cases, which have been registered using our approach. Despite the difficulty with validating non-rigid registration, particularly for multimodal applications, we have been able to demonstrate that the error due to registration (compared against a clinical evaluation) is typically less than 20% of the local displacement resulting from elastic deformation. This is seen as an encouraging result.

Our ongoing effort in this research is aimed at developing a more robust registration framework in the context of a larger and (as indicated in the discussion) more diverse validation study. We have implemented the data fusion framework as a software package, which is currently installed at the John Radcliffe Hospital, Oxford, where it is used on a regular clinical basis. From our clinical experience with our software tools, we have found that there is a considerable gulf between a theoretical fusion concept with the implication of automatic performance, and a clinically useable (and accepted) system. Our long-term objective is to produce a system capable of enabling the radiologist to do a routine comparison between mammography and MRI

Acknowledgements

The authors would like to thank many people for their contribution to this work. This includes Dr. Jérôme Declerck (Mirada Solutions) for useful discussions about this work, Mr. Dermot Dobson and the staff at the Oxford MRI Centre (John Radcliffe Hospital), Dr. Eve Coste-Manière, Dr. Hervé Delingette, Professor Nicholas Ayache and Dr. Alexis Roche (Mirada Solutions Ltd.) for useful feedback and discussion. The idea of using MRI to potentially validate breast tissue deformation was the inspiration of Professor Elliot McVeighe (Johns Hopkins/NIH Bethesda) as well as the preliminary breast data. Finally, C. Behrenbruch would like to gratefully acknowledge the Commonwealth Scholarship and Fellowship Program for his doctoral funding at Oxford University. P. Armitage and J.M. Brady acknowledge the EPSRC for funding (GR/M54995).

References

Arad, N., Reisfeld, D., 1995. Image warping using few anchor points and radial functions. Computer Graphics Forum 14 (1), 35–36.

Armitage, P., Behrenbruch, C.P., Brady, M., 2001. A fast, reliable method for calculated T_1 in 3D contrast-enhanced breast MR imaging. In: Proc. ISMRM'01, Glasgow.

Azar, F.S., Metaxas, D.N., Schnall, M.D., 2000. A finite element model of the breast for predicting mechanical deformation during biopsy procedures. In: Proceedings of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis.

Behrenbruch, C.P., Marias, K., Armitage, P.A., Brady, J.M., Clarke, J., Moore, N., 2000a. The generation of simulated mammograms from contrast-enhanced MRI for surgical planning and postoperative assessment. In: Proceedings of IWDM (the International Workshop in Digital Mammography), Toronto, Canada. Kluwer Academic Publishers, Dordrecht.

Behrenbruch, C.P., Marias, K., Armitage, P.A., Yam, M., Moore, A., English, R.E., Brady, J.M., 2000b. Mammography-MRI 2D/3D data fusion for breast pathology assessment. In: Proceedings of MICCAI (Medical Robotics, Imaging and Computer Assisted Intervention). Springer, Pittsburgh, PA.

Besl, P., McKay, N., 1992. A method for registration of 3-D shapes. IEEE Trans. PAMI 14, 239–256.

Bookstein, F.L., 1989. Principal warps: ''thin-plate splines and the decomposition of deformations''. IEEE Transactions on Pattern Analysis and Machine Intelligence 11 (6), 567–585.

Bovik, A.C., Clark, M., Giesler, W.S., 1990. Multichannel texture analysis using localized spatial filters. IEEE PAMI 12, 55–73.

Brix, G., Semmler, W., Schad, L.R., Layer, G., Lorentz, W., 1991. Pharmacokinetic parameters in CNS Gd-DTPA enhanced MR imaging. J. Comput. Assist. Tomogr. 15, 621–628.

Buckley, D.L., Kerslake, R.W., Blackband, J., Horsman, A., 1994. Quantitative analysis of multi-slice Gd-DTPA enhancement dynamic MR images using an automated simplex minimisation procedure. Magnetic Resonance in Medicine 32, 646–651.

Burt, P.J., Adelson, E.H., 1983. The Laplacian pyramid as a compact image code. IEEE Transactions on Communication 31 (4), 540–552.

Christensen, G.E., Rabbitt, R.D., Miller, M.I., 1994. 3D brain mapping using deformable neuroanatomy. Physics in Medicine and Biology 39, 609–618.

Coifman, R.R., Wickerhauser, M.V., 1992. Wavelets and adapted waveform analysis. In: Benedetto, Frazier, (Eds.), Wavelets: Mathematics and Applications. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL.

Coifman, R.R., Meyer, Y., Wickerhauser, M.V., 1991. Adapted waveform analysis, wavelet-packets and applications. In: Proceedings of ICIAM'91. SIAM Press, Philadelphia, PA, pp. 41–50.

Coifman, R.R., Meyer, Y., Quake, S.R., Wickerhauser, M.V., 1992. Signal processing and compression with wavelet packets. In: Meyer, Y., Roques, S. (Eds.), Proceedings of the International Conference 'Wavelets and Applications', Toulouse, France, 8–13 June, Observatoire Midi-Pyrénées de l'Université Paul Sebatier, pp. 77–93.

Cotin, S., Delingette, H., Ayache, N., 1999. Real-time elastic deformation of soft tissues for surgery simulation. IEEE Transactions on Visualization and Computer Graphics 5 (1), 62–73.

D'Astous, F., Jernigan, M.E., 1984. Texture discrimination based on detailed measures of the power spectrum. In: Proceedings of the 7th International Conference on Pattern Recognition, pp. 83–86.

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

Daubechies, I., Grossman, A., Meyer, Y., 1986. Painless nonorthogonal expansions. Journal of Mathematical Physics 27, 1271–1283.

Declerck J., Feldmar J., Goris M.L., Betting F., 1996. Automatic registration and alignment on a template of cardiac stress and rest SPECT images. Mathematical Methods in Biomedical Image Analysis.

Evans, A.C., Dai, W., Collins, L., Neelin, P., Marrett, S., 1991. Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis. In: Leow, X. (Ed.), Proc. SPIE 1445, Medical Imaging V: Image Processing. San Jose, CA, USA, pp. 236–246.

Fletcher, S.W., Black, W., Harris, R., Riner, B.K., Shapiro, S., 1993. Report of the International Workshop on Screening for Breast Cancer. J. Natl. Cancer Inst. 85, 1644–1656.

Gilhuijs, K.G.A., 1995. In: Automated Verification of Radiation Treatment Geometry. Universiteit van Amsterdam, Febo-druk, Enschede-Utrecht, The Netherlands.

Grossman, A., Morlet, J., 1985. Decomposition of functions into wavelets of constant shape and related transforms. In: Streit, L. (Ed.), Mathematics and Physics, Lectures on Recent Results. World Scientific, Singapore.

Haralick, R.M., Shanmugam, K., Dinstein, I., 1973. Texture features for image classification. IEEE SMC 3 (6), 610–621.

Hardy, R.L., 1981. Multiquadric equations of topography and other irregular surfaces. Journal of Geophysical Research 76 (8), 1905–1915.

Hayton, P., Brady, M., Tarassenko, L., Moore, N., 1997. Analysis of dynamic MR breast images using a model of contrast enhancement. Medical Image Analysis 1 (3), 207–224.

Highnam, R.P., Brady, J.M., 1999. In: Mammographic Image Processing. Kluwer Academic Press, Dordrecht.

Holschneider, M., 1995. In: Wavelets: An Analysis Tool. Oxford Mathematical Monographs. Oxford University Press, London.

Horn, B.K.P., Schunck, B.G., 1981. Determining optical flow. Artificial Intelligence 17, 185–203.

Karssemeijer, N., Thijssen, M., Hendriks, J., van Erning, L. (Eds.), 1998. Digital Mammography, Computational Imaging and Vision. Kluwer Academic Publishers, Dordrecht.

Kavanagh, A.M., Mitchell, H., Giles, G.G., 2000. Hormone replacement therapy and accuracy of mammographic screening. Lancet 355, 270–274.

Kita, Y., Highnam, R., Brady, M., 1998. Correspondence between different view breast X-rays using a simulation of breast deformation. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 700–707.

Koenderink, J.J., 1984. The structure of images. Biological Cybernetics 50, 363–370.

Kovesi, P., 1996. Invariant measures of imag4e features from phase information. Ph.D. thesis, Department of Psychology, University of Western Australia.

Kyriacou, S.K., Davatzikos, C., Zinreich, S.J., Bryan, R.N., 1999. Nonlinear elastic registration of brain images with tumour pathology using a biomechanical model. IEEE Trans Med. Imag. 18 (7), 580–592.

Lanyi, M., 1996. In: Diagnosis and Differential Diagnosis of Breast Calcifications. Springer, Berlin.

Leung, W., Goldberg, F., Zee, B., Sterns, E., 1997. Mammographic density in women on postmenopausal hormone replacement therapy. Surgery 122 (4), 669–673 (discussion 673–674).

Litherland, J.C., Evans, A.J., Wilson, A.R., 1997. The effect of hormone replacement therapy on recall rate in the National Health Service Breast Screening Programme. Clin Radiol 52 (4), 276–279.

Little, J.A., Hill, D.L.G., Hawkes, D.J., 1997. Deformations incorporating rigid structures. Computer Vision and Image Understanding 66 (2), 223–232.

Mallat, S., 1998. In: A Wavelet Tour of Signal Processing. Academic Press, San Diego, CA.

Mallat, S.G., 1989. Multiresolution approximation and wavelet orthonormal basis of $L^2(R)$. Transactions of the AMS 315, 69–87.

Marias, K., Brady, J.M., Parbhoo, S., Seifalian, A.M., 1999. Registration and matching of temporal mammograms for detecting abnormalities. In: Proceedings of Medical Imaging Understanding and Analysis (MIUA'99). Oxford, UK.

Marias, K., Behrenbruch, C.P., Brady, J.M., Parbhoo, S., Seifalian, A.M., 2000. Multi-scale landmark selection for improved registration of temporal mammograms. In: Proceedings of IWDM (International Workshop in Digital Mammography), Toronto, Canada. Elsevier, Amsterdam.

Marr, D., 1982. In: Vision: A Computational Investigation Into the Human Representation and Processing of Visual Information. W.H. Freeman, San Francisco, CA.

Nystrom, L., Rutqvist, L.E., Wall, S., Lindgren, A., Lindqvist, M., Ryden, S., Andersson, I., Bjurstam, N., Fagerberg, G., Frisell, J., Tabar, L., Larsson, L.G., 1993. Breast cancer screening with mammography: overview of Swedish randomised trials. Lancet 341, 973–978.

Picinbono, G., Delingette, H., Ayache, N., 2000. Real-time large displacement elasticity for surgery simulation: non-linear tensor-mass model. In: Proceedings of MICCAI (Medical Robotics, Imaging and Computer Assisted Intervention), Pittsburgh, USA. Springer, Berlin.

Pincus, S.M., 1992. Approximate entropy: statistical properties and applications. Commun. Statist. Theory Methods 21, 3061–3077.

Pincus, S.M., 1997. Not all (possibly) 'random' sequences are created equal. Proc. Natl. Acad. Sci. USA 94, 3513–3518.

Rankin, S.C., 2000. MRI of the breast (Review Article). Br. J. Radiol. 73, 806–818.

Reuckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J., 1999. Nonrigid registration using free-form deformations: application to breast MR images. IEEE Transactions in Medical Imaging 18 (8), 12–21.

Rohr, K., Stiehl, H.S., Sprengel, R., Beil, W., Buzug, T.M., Weese, J., Kuhn, M.H., 1996. Point-based elastic registration of medical image data using approximating thin-plate splines. Lecture Notes in Computer Science 1131, 297–306, Springer.

Ruprecht, D., Müller, H., 1993. Free form deformation with scattered data interpolation methods. Computing Supplementum 8, 267–281.

Sallam, M.Y., Bowyer, K.W., 1996. Registration and difference analysis of corresponding mammogram images. Medical Image Analysis 1 (1), 73–91.

Stein, E., 1979. In: Singular Integrals and the Differentiability Property of Functions. Princeton University Press, Princeton NJ.

Sweldens, W., 1996. The lifting scheme: a custom-design construction of biorthogonal wavelets. Appl. Comput. Harmon. Anal. 3 (2), 186–200.

Thirion, J.-P., 1998. Image matching as a diffusion process: an analogy with Maxwell's demons. Medical Image Analysis 2 (3), 243–260.

Tofts, P.A., Kermode, A., 1991. Measurement of the blood–brain barrier and leakage space using dynamic MR imaging 1. Fundamental concepts. Magnetic Resonance in Medicine 17, 347–367.

Turner, M.R., 1986. Texture discrimination by Gabor functions. Biological Cybernetics 55, 71–82.

Verbeek, A.L.M., Hendricks, J.H.C.L., Holland, R., Mravunac, M., Strumans, F., Day, N.E., 1984. Reduction in breast cancer mortality through mass screening with modern mammography. (First results of the Nijmegen Project 1975–1981). Lancet i, 1222–1224.

Warren, R.M.L., Hayes, C., Advisory Group of the UK Study of MRI Screening for Breast Cancer – MARIBS, 2000. Localization of breast lesions shown only on MRI – a review for the UK study of MRI screening for breast cancer. Br. J. Radiology 73, 123–132.

Wickerhauser, M.V., 1994. In: Peters, A.K. (Ed.), Adapted Wavelet Analysis from Theory to Software. IEEE Press.

Wirth, M.A., Choi, C., 2000. The comparative analysis of mammograms: a nonrigid approach to registration. In: Proceedings of IWDM (International Workshop in Digital Mammography), Toronto, Canada. Elsevier, Amsterdam.

Witkin, A.P., 1983. Scale-space filtering. In: 8th International Joint Conference on Artificial Intelligence. Karlsruhe, Germany.

Wolfe, J.N., 1982. In: Xenoradiography of the Breast. Thomas, Springfield, IL.

Yam, M., Brady, J.M., Highnam, R.P., English, R.E., Kita, Y., 1999. Reconstructing microcalcification clusters in 3-D using a parameterized breast compression model. In: Computer Assisted Radiology and Surgery (CARS), Paris.

Zhu, S.C., Yuille, A., 1996. Region competition: unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. IEEE Trans. PAMI 18 (9), 884–900.