

## A non-rigid registration algorithm for dynamic breast MR images

Paul M. Hayton<sup>a,\*</sup>, Michael Brady<sup>a</sup>, Stephen M. Smith<sup>b</sup>, Niall Moore<sup>c</sup>

<sup>a</sup> *Department of Engineering Science, Oxford University, Oxford OX1 3PJ, UK*

<sup>b</sup> *Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB),  
Department of Clinical Neurology, Oxford University, Oxford, UK*

<sup>c</sup> *MRI Centre, The John Radcliffe Hospital, Oxford, UK*

Received 19 April 1999; received in revised form 28 June 1999

---

### Abstract

Magnetic resonance image analysis is a promising technique for diagnosing breast cancer, particularly in women for whom X-ray mammography is ineffective. If breast motion is not corrected for, diagnostic accuracy is significantly reduced. In this paper, we analyse the kinds of motion that arise during image formation and we describe a model based non-rigid registration algorithm to estimate and correct for breast motion. Registration of breast MR images is complicated by the use of a contrast agent which results in a non-uniform increase in intensity across the image. The work described here forms part of an implemented breast MR analysis system which allows automatic detection and segmentation of regions of focal enhancement and non-rigid image registration. © 1999 Elsevier Science B.V. All rights reserved.

**Keywords:** Optic flow; 2D motion estimation; Correlation; Uncertainty; Contrast-enhanced MRI

---

### 1. Introduction

Breast cancer is a major problem for public health in the Western world, where it is the most common cancer among women. In the European Community, for example, breast cancer represents 19% of cancer deaths and fully 24% of all cancer cases. It is diagnosed in a total of 348,000 cases annually in the USA and EC [8] and kills almost 115,000 annually. Approximately 10% of women will develop breast cancer during the course of their lives.

In almost every field of medical imaging, clinicians and technicians rely increasingly on computer processing of images as a key part of their decision making. This can involve

---

\* Corresponding author. Email: pmh@robots.ox.ac.uk.

framing a diagnosis or deciding when to call for further investigation such as a repeat image, an image of a different type, or a biopsy. Images are also used to monitor therapy. Unfortunately, as applications of image analysis go, medical applications are tough in general, and breast cancer image analysis is one of the toughest. There are many reasons for this: intrinsically complex and noisy images; the highly variable and irregular shapes of the objects of interest; abnormalities that appear as quite subtle, irregular, non-local differences from their surrounds; changing imaging conditions; and patient motions. Add to this the increasing need for quantitative information, precision, and reliability (very few false positives), and the image processing challenge becomes quite daunting.

The combination of task complexity and need for accuracy in the face of ubiquitous uncertainty implies the need for image analysis programs to mobilise knowledge about image formation and about the appearances of normal and abnormal tissue to be mobilised by image processing algorithms. This paper describes a key part of a knowledge-based system that we have implemented at the John Radcliffe Hospital Trust in Oxford to support the diagnosis of breast cancer using magnetic resonance imaging (MRI). Chapter 14 of [14] describes the system in more detail and presents some retrospective clinical trials.

In the remainder of the Introduction, we motivate our reasons for developing contrast-enhanced MRI analysis of breast cancer and we show why estimating the motion of a patient's breasts during the 7 minute acquisition period is critical. We then analyse the kinds of motions that are likely to occur during the examination, since these define the type of registration algorithm that has to be developed. It turns out that it is necessary to develop a non-rigid registration algorithm that can operate reliably despite the fact that the image intensities are changing rapidly and non-uniformly throughout the breast. We discuss a number of local image similarity measures and show why mutual information between two images is well suited to the current task. With this background, we develop a Bayesian non-rigid registration algorithm in Section 6. We describe an implementation that operates at several spatial scales in Section 7, then show why it is necessary to compensate for an image degradation factor called the "bias field" before presenting a set of results in Section 9. A fuller account of our results from a clinical trial have been presented elsewhere.

### *1.1. Mammography and the need for contrast-enhanced breast MRI*

Current image-based diagnosis and screening for breast cancer is based predominantly on (X-ray) mammography and it has enjoyed considerable success, leading to a reduction in the mortality rate, particularly in countries such as the UK where there is a national screening program. X-ray mammography is currently the only viable mass screening modality due its effectiveness and low cost. The UK breast screening program, for example, currently checks annually over one and a half million women between the ages of 50 and 65 and repeats the examination every three years. The reason for restricting to this age range is the following: when a woman reaches the menopause, milk bearing glandular breast tissue, which was previously opaque to X-rays, starts to turn to fat, which is relatively transparent to X-rays, and so abnormalities, characterized by denser tissue, are visible as brighter areas in a mammogram.

Despite the diagnostic success based on X-ray mammography there is room for considerable improvement: currently more than one operation in two is performed to remove a lesion that turns out to be benign; 8–25% of cancers are missed, and 70–80% biopsies turn out to be benign. So, although mammography is highly sensitive, findings are frequently non-specific (Ciatto et al. [4]; Feig [9]; Ikeda and Anderson [17], Jackson et al. [18]; Sickles [25]). For example, although some radiological findings are characteristic of benign or malignant disease (e.g., teacup microcalcifications and spiculated opacity, respectively), others are quite indeterminate. Also, certain fibrocystic breast diseases have a high risk of cancer, but screening is difficult because often women present at examination with dense breasts, in which it is difficult to detect abnormal opacities.

One reason that X-ray mammograms yield indifferent diagnoses is the poor quality of the images. This has led to considerable effort in mammographic image processing (see, for example, Highnam and Brady [14]. Such work will, and must, continue to develop since X-ray mammography is likely to remain the principle imaging modality for the next few years. However, X-ray mammography, as it is currently understood, inevitably has a number of fundamental shortcomings:

- *Projective*: to reduce patient dosage, the breasts are compressed (tightly) between parallel plates while a mammogram is taken. It follows that localisation of a lesion or microcalcification on the basis of an X-ray mammogram is inherently poorer than might be hoped for with truly three-dimensional imaging.
- *Limited applicability*: Mammography is effective only for a limited class of women. It is not only younger women that are excluded (for the reasons outlined above): X-ray mammography is also ineffective for women with scar tissue following surgery, some women undergoing Hormone Replacement Therapy (HRT), women with mamographically occult tumours and women whose tumour is only visible in one view.
- *Dynamics of tumours*: Mammograms present an image of breast anatomy. However, malignant cancers are characterised more by their dynamic behaviour. To nourish its growth, a tumour sprouts a network of new blood vessels (neovascularisation) and this causes a sharp increase in the local blood supply (angiogenesis), which mammography cannot image.

The upshot of these limitations is that imaging techniques complementary to X-ray mammography have been developed over the past decade. Of particular interest are ultrasound (2D or 3D, various kinds of Doppler imaging), nuclear medicine (PET and Spect), and contrast-enhanced MRI (see Chapters 14 and 15 of Highnam and Brady [14] for introductions to these). Each has advantages and disadvantages, and it is currently unclear which will eventually predominate or whether two or more will be used in combination.

Breast MRI can overcome the limitations listed above: the breast is imaged in 3-D (as a set of 2-D slices); it is effective for all the classes of women for which X-ray mammography is ineffective, not least younger women; and it can give dynamic information about angiogenesis. However, a disadvantage of *standard* MRI is that it does not work nearly as well as X-ray mammography in discriminating between abnormalities and their surrounding glandular tissue. Indeed, as early as 1992, Heywang-Köbrunner [13] summarised research findings by concluding that apart from the diagnosis of fibroadenomas there is no widely accepted role for plain MR imaging of the breast.

For this reason, the identification of abnormalities currently requires a contrast agent to be injected into the patient's bloodstream. To date, the only contrast agent that has as yet been approved by the Food and Drugs Administration (FDA) for clinical use is the paramagnetic chelate salt Gd-DTPA of Gadolinium.

In a typical examination, a patient is scanned prior to and immediately after intravenous injection of the contrast agent. Post-Contrast scans are performed repeatedly for approximately 7 minutes producing a total of 8 scans. The MR image consists of a set of 2D images of parallel tissue slices that may not be contiguous. In our (typical) case, a volume consists of 16 slices, each approximately 5 mm thick and comprising  $256 \times 256$  pixels, each pixel corresponding to a tissue volume measuring approximately  $1 \text{ mm} \times 1 \text{ mm}$  in the image plane. There is a gap of 3 mm between each slice. It follows that the volume imaged measures  $256 \times 256 \times 125 \text{ mm}^3$ . We use a dedicated double breast coil and the volume imaged covers the whole of both breasts.

Because of the increased angiogenesis, the intensities of image elements surrounding a tumour increase rapidly and by a large amount, while those corresponding predominantly to fat (which contains relatively small amounts of blood) hardly change. Normal tissue voxels rise, but more slowly and less than voxels that are abnormal. The analysis of the rise in the MR signal level when Gd-DTPA is injected into the blood-stream is key to the diagnosis of a tumour and the shape of this enhancement provides indications of possible malignancy.

### *1.2. The need for registration in contrast enhanced breast MRI*

Since a breast MR examination involves scanning a patient before and repeatedly after injection of the contrast agent, if there is any movement of the patient's breasts during the course of the examination this enhancement may be distorted or undetected. This motion could be due simply to the breathing of the patient, or due to other movement made during the relatively long (7 minutes) MRI examination. In order to accurately analyse the enhancement of an abnormality, and to delineate the shape of the tumour as precisely as possible, this motion must be compensated for. Several studies on contrast-enhanced breast MRI have reported problem cases of patients who have moved during the scan (e.g., Porter [22]). Pretorius et al. [23] concluded that in order for any quantitative analysis to be reliable the problem of registration must first be addressed. The following sections review previous work on registration and introduce a novel registration algorithm that enables breast motions to be estimated and corrected for.

## **2. Types and causes of motion**

There are several potential causes of motion during breast MR examinations. In order to devise a suitable registration algorithm, we need to first fully understand the problem, namely the type and extent of these motions.

As is shown in Fig. 2, during an MR examination, the patient lies prone, normally with her arms forward. The patient's body is supported on the couch, so assuming the patient is not deliberately trying to move, translation in the *Z* direction is less likely than motion




Fig. 1. Example images showing movement due to relaxation of the pectoral muscle. The breast is predominately fat (high signal intensity) and the pectoral muscle is the darker grey tissue towards the top of the image. In the post-contrast image, the patient has relaxed her pectoral muscle, changing its shape and causing non-rigid movement in the surrounding tissue. The result of applying the registration algorithm to these images will be shown later (Fig. 8).




Fig. 2. The patient lies prone inside the MR machine. The axial slice plane ( $XY$  plane), which is the imaging plane for the breast MR images in this thesis, is indicated.

in the  $XY$  plane (which is the axial imaging plane) since it requires considerable effort to move forwards or backwards whilst lying on one's front. The effect of any translation in the  $Z$  direction is further reduced by the 5 mm slice thickness—i.e., a 1 mm motion in the  $Z$  direction only results in a 20% partial volume change.

Large movements do occasionally occur when patients are uncomfortable during the scan. Due to the patient lying prone these movements tend to be rotational about the  $Z$ -axis as the patient changes position on the couch. Rotations about the  $X$ -axis are unlikely for the same reason as translations along the  $Z$ -axis, namely the patient is lying on her




Fig. 3. A typical MR image of both breasts (left) and a diagrammatic representation of the position of the double breast coil. The breasts are constrained within the double breast coil, so an upper bound on size of motion of the breasts can be estimated.

front. Once again, because the slice thickness is 5 mm, the effects of small rotations about the  $X$  and  $Y$  axis will be reduced.

We can put a maximum bound on the movement of the patient's breasts since they hang pendulous inside the double breast coil (Fig. 3). Motion is constrained; the breasts will remain inside the double breast coil. Therefore, for an average woman, an upper bound on the expected translation or displacement of a point can be estimated at about 25 mm.

In addition to any explicit patient movement, tensing or relaxing of the pectoral muscles manifests itself as movement which affects the whole of the breast via the 'Coopers ligaments' which run throughout the breast. In Fig. 1 the patient's pectoral muscles were initially tense during the pre-contrast scan. After the pre-contrast scan, she relaxed and the shape of the pectoral muscle changed causing movement of the surrounding tissue.

### 3. Review of non-rigid registration techniques

For dynamic breast MR image sequences, it is non-rigid registration techniques that are of most interest since the human breast moves non-rigidly. In the case of the work presented here, the problem can only be considered in 2D currently due to the sparse nature of the MR scans ( $1 \times 1 \times 5$  mm voxels with 3 mm gaps between slices), however, as we argued in the previous section, most movement is within the axial slice plane. At the end of this paper, the possibility of performing a complete 3D non-rigid registration of breast MR images is discussed.

Maintz et al. [19] provide a detailed survey of medical image registration techniques. For breast MR images, only intrinsic registration algorithms are useful. These differ from each other firstly in the information from the images that is used by the algorithm, and secondly by the method for finding a globally optimum solution. The information used to register images can be any object that is identifiable in both images as corresponding to the same anatomical location. Algorithms can be separated into two categories by the type of information used:

- (1) *Landmark and segmentation based algorithms* (Sparse matching): Geometric features or points which hopefully correspond to anatomical structures are extracted from each image. These are matched between the images so that a transformation can be fitted. This is a form of information reduction where the image is reduced to a set of points, contours or surfaces from which the motion can be estimated. There are, however, few reliable geometric features in breast MR images.
- (2) *Voxel property based algorithms* (Dense matching): These use the intensities of the voxels themselves and define a measure of similarity between different images/sub-images. The registration parameters are optimized to maximize this measure. It is into this category that the algorithm proposed in this paper falls.

Segmentation-based methods in general have proved very successful in registering images of the brain. Brain images are, however, ideally<sup>1</sup> piecewise constant, so there are many features that can reliably be extracted. In breast MR images, there are few geometric features other than the breast edge that can in general be detected automatically. (There are other features but they cannot be relied upon.) Sufficient information is required to estimate the motion of the breast, which in some cases is quite complex, so voxel intensity based algorithms are of more interest for this application.

### 3.1. Voxel-property based algorithms

Voxel-property based algorithms use the full image content to determine the transformation between two images. These algorithms require a measure of similarity to be defined between the images or sub-images. The differences between these algorithms are the choice of similarity measure, determined by the modality and characteristics of the images, and the model of deformation allowed, determined by the object being imaged. Optical flow algorithms, applied to medical images, also fall into this category.

In [31], Vemuri proposes a novel algorithm to identify the motion field between two images which incorporates a modified Newton method into a spline-based optical flow framework. The algorithm uses the sum of squared differences as its measure of similarity. The innovative contribution of the work is the minimization framework which pre-computes the Hessian matrix of the motion field at the optimum prior to using it at each iteration. However, the method assumes the Hessian is constant and that the location of the optimum is known approximately. The motion field, and implicitly the deformation, is represented in a B-spline basis.

Bro-Nielsen [2] used a fluid flow model for the deformation. His work is based on that of Christensen [3] who developed the viscous fluid flow model. This computes a force on the transformation composed of an image force, which alters the transformation such that the images move toward one another, and a viscous fluid force derived from a partial differential equation describing fluid flow. The transformation is adjusted at each iteration in the direction of this force until an equilibrium stopping condition is reached. Thirion [30] performs non-rigid registration using what he terms ‘demons’ which drive the image transformation based on the difference in image intensity gradients. Bro-Nielsen [2]

---

<sup>1</sup> Although brain MR images are ideally piecewise constant, partial volume effects and the bias field degrade the images.

showed that ‘demons’ are very similar to the fluid flow algorithm, and can be considered as a special case of that algorithm.

Viola [32] proposed a algorithm to register images by maximization of mutual information. Mutual information, measures the extent to which two images are functionally related to one another. The algorithm seeks to optimize the parameters of a transformation (rigid or affine) to maximize the mutual information between the two images. Wells and Viola [34] applied this algorithm to register an MR volume with CT and PET volumes. The complexity of the algorithm is dependent on the number of parameters of the transformation. Mutual information will be discussed below.

#### 4. Optic flow methods

Optical flow based methods derive their similarity measure from the *motion constraint equation*, proposed by Horn and Schunck [16] which states:

$$\nabla I \cdot \frac{d\mathbf{r}}{dt} = -\frac{\partial I}{\partial t}, \quad \text{where } \mathbf{r} = \begin{pmatrix} x \\ y \end{pmatrix},$$

where  $I(x, y, t)$  is the image. This is a linear spatial and temporal approximation and implies that brightness is conservative, i.e., we expect pixels in two different images which correspond to the same physical point in space to have the same intensity. Most algorithms for determining optic flow do so by minimizing an energy function based on this equation.

A non-rigid transformation or motion field is usually computed by minimizing an energy function which combines a similarity measure with the model of deformation. This energy function is typically of the form:

$$E(\mathbf{u}) = \int_A S(\mathbf{r}, \mathbf{u}) + \text{Reg}(\mathbf{u}) dA, \quad \text{where } \mathbf{u} = \begin{pmatrix} u(x, y) \\ v(x, y) \end{pmatrix}, \quad (1)$$

where  $A$  is the region of interest,  $S(\mathbf{r}, \mathbf{u})$  is the measure of similarity between the two images at position  $\mathbf{r}$ , with a displacement  $\mathbf{u}(x, y)$  between the images and  $\text{Reg}(\mathbf{u})$  is a regulariser. In minimizing this energy function, a solution is sought whereby under the transformation, the images are similar subject to the similarity measure whilst the motion field conforms to the deformation model (for optic flow, this is  $\text{Reg}(\mathbf{u})$  in the equation above which is typically of a form similar to  $\nabla^2 \mathbf{u}$  which imposes second order smoothness).

Many existing methods work on the principle of finding an optimal value for the motion at each pixel and then reducing errors by some sort of smoothing of the resulting vector field. These errors arise largely because of two factors. The first of these is flow measurement inaccuracy, arising from the various noise processes which occur throughout the image generation, along with non-modelled changes in the image. The second cause of “error” is the well known aperture effect [16] such that from the motion constraint equation, only the component of  $\mathbf{u}$  perpendicular to the direction of  $\nabla I$  (i.e., motion perpendicular to an edge) can be estimated directly. The component of motion parallel to the edge must be inferred from measurements elsewhere.



Various methods of overcoming these problems have been proposed. 2D-feature-based optic flow methods (e.g., [27,28]) overcome the aperture problem by calculating flow only at places in the image where the 2D flow field is well-conditioned. 1D-feature-based (edge-based) methods (e.g., [10,15]) find the normal flow at edges, and interpolate around contours to enable estimation of the flow component parallel to the edges. However, none of these methods carry out very much error suppression on the flow estimates. The simplest gradient-based methods (e.g., [16]) reduce flow error by effectively applying linear smoothing to the flow field, whilst more advanced methods (e.g., [20,21]) apply anisotropic smoothing, and try not to smooth flow across image edges or motion boundaries.

A weakness with these approaches is that no information related to possible alternatives to the “optimal” flow vector is considered; thus much potentially useful information is lost. When a first estimate of the motion field is made from the similarity measures,  $u_0(\mathbf{r})|S(\mathbf{r}, \mathbf{u})$  a choice is forced and the algorithm then spends time *unpicking* the errors computing  $u_i(\mathbf{r})|u_{i-1}(\mathbf{r})$  at each iteration to give a final estimate  $u_{\text{final}}(\mathbf{r})$ . There can be several plausible motions locally and algorithms often fall into local minima; They have lost/ignored the initial evidence that existed for the final solution i.e.,  $u_{\text{final}}(\mathbf{r})|S(\mathbf{r}, \mathbf{u})$ .

The reader is referred to [1] for an excellent review of optic flow methods and a related approach for applying Bayesian methods to optical flow estimation can be found in [26].

Computer-based registration algorithms are forced to analyse images at a low level. As stated in the previous section, in developing or choosing a strategy for registration of breast MR images, we are faced with two problems:

- Finding a similarity measure suitable for breast MR images which is robust under contrast enhancement and motion artifacts.
- Developing/selecting a method for combining local similarity information under a model of deformation to obtain a global transformation or motion field between the images.

The following sections discuss the choice of similarity measure for breast MR images then introduce a novel algorithm for finding a solution from the similarity data. The model of deformation that is used is that the optimum global deformation is that which minimizes the Laplacian of the motion field  $\nabla^2 \mathbf{u}$ . This is a simplification of the linear elastic material model [2]. The assumption of linear elasticity is a reasonable first model for the breast. Other models are discussed later in the conclusion.

## 5. Voxel similarity measures

Any voxel property based registration algorithm requires a method of determining whether two subimages represent the same anatomical location. For breast MR images this must be considered carefully. MR signal intensities are subject to change between images due to contrast enhancement—this change can be of the order of 100%. Indeed, it is this non-uniform enhancement that is the key to diagnosis.

The measure of similarity is often implicit in the final registration algorithm. In the classic optic flow algorithm by Horn and Schunck [16], the measure of similarity is part of the energy function to be minimized (as in Eq. (1)) from which an update equation is derived in terms of image derivatives. To avoid the use of noisy image derivatives,

the algorithm presented in this chapter calculates the similarities between different image locations explicitly, i.e., the similarity measure is some function of the signal intensities at two image locations:

$$S(\mathbf{r}, \mathbf{u}) = f(I(\mathbf{r}), I(\mathbf{r} + \mathbf{u})), \quad \text{where } \mathbf{r} = \begin{pmatrix} x \\ y \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} u(x, y) \\ v(x, y) \end{pmatrix}.$$

Bro-Nielsen [2] gives a survey of a wide range of voxel similarity measures which he relates to texture analysis in terms of *Grey Level Co-occurrence Matrices*. The most interesting of these similarity measures for breast MR registration are discussed here.

### 5.1. Correlation measures

The simplest method for comparing two image locations is the error norm, applied to individual pixels or over small regions.

$$E(\mathbf{r}, \mathbf{u}) = \frac{1}{N} \sum_{ROI} |I_1(\mathbf{r}) - I_2(\mathbf{r} + \mathbf{u})|^n,$$

$$S(\mathbf{r}, \mathbf{u}) = \frac{1}{1.0 + E(\mathbf{r}, \mathbf{u})},$$

where  $N$  is the number of pixels in the ROIs centred on  $\mathbf{r}$  in image 1 and  $\mathbf{r} + \mathbf{u}$  in image 2. For  $n = 1$ , the  $L_1$  norm, this is the absolute intensity difference; for  $n = 2$ , the  $L_2$  norm, it is the square Euclidean distance. A more appropriate correlation measure for breast MRI, however, is normalized “centered” cross-correlation.

$$S(\mathbf{r}, \mathbf{u}) = \frac{\sum (I_1(\mathbf{r}) - \mu_1)(I_2(\mathbf{r} + \mathbf{u}) - \mu_2)}{(\sum (I_1(\mathbf{r}) - \mu_1)^2 \sum (I_2(\mathbf{r} + \mathbf{u}) - \mu_2)^2)^{1/2}},$$

$$\text{where } \mu = \frac{1}{N} \sum_{i=0}^{N-1} I(\mathbf{r}_i).$$

Devernay [7] relates normalized cross-correlation to the cosine rule applied to the centered intensity vectors  $((I_1(r_0) - \mu_1), (I_1(r_1) - \mu_1), \dots)^T$  and  $((I_2(r_0 + \mathbf{u}) - \mu_2), (I_2(r_1 + \mathbf{u}) - \mu_2), \dots)^T$  where normalized cross-correlation is the cosine of the angle between these vectors. It is also a measure of the least squares residual error from fitting a line to a set of data and consequently for registration can be expected to perform best when there exists a linear relationship between the image intensities.

### 5.2. Mutual information

Mutual information [32] (or relative entropy) is a statistical measure that describes how likely it is that a random variable  $Y$  is functionally related to  $X$ . The *mutual information* is defined in terms of entropy  $H$  as:

$$M(X, Y) = H(Y) - H(Y|X).$$

The conditional entropy can be expressed in terms of the joint entropy:

$$H(Y|X) = H(X, Y) - H(X)$$

so that the mutual information can be written as

$$M(X, Y) = H(X) + H(Y) - H(X, Y), \quad (2)$$

which under a transformation  $T$ , for probability density estimates  $a(x)$  and  $b(x)$ , becomes:

$$M(a(x), b(T(x))) = H(a(x)) + H(b(T(x))) - H(a(x), b(T(x))).$$

In maximizing this measure, transformations are encouraged which project onto the complex parts of  $b$  such that  $a$  and  $b$  are functionally related.

Mutual information is consequently a statistical measure for which the only assumption required is that two pixels that have the same intensity in the first image also have the same intensity in the second image. In Breast MR images, this is not always true over the whole image—places where it does not hold are of diagnostic interest. It is, however, likely to be true on a local scale over most of the image.

Variations of mutual information have been proposed for registration of medical images. Studholme [29] proposed the measure of *normalized mutual information*, which is defined as the ratio of joint and marginal entropies of the images and is designed to consider the amount of mutual information with respect to the information provided by the individual images. Without this normalization, a registration algorithm will favour transformations which result in high marginal entropies of the two images, rather than transformations which maximize the *proportion* of information shared between the images. The only problem with this normalization is that it results in two image locations which are similar but have very little information (i.e., uniform intensity) being considered as important as two locations which are rich in information (i.e., textured).

### 5.3. Comparing similarity measures

An experiment was performed to compare the similarity measures under contrast enhancement. Two artificial images were created; Image 1 has a single *feature* (to simulate an enhancing blood vessel or abnormality). Image 2 was created from image 1 by translating it by 5 pixels, rotating it by  $10^\circ$  and increasing the intensity of the feature by 120%. The similarity measures were tested by selecting a point for the similarity kernel in image 2 and translating the kernel in image 1 by an amount  $-20 < dx < 20$  and calculating the similarity measures between the two kernels. The correct result is a maximum similarity at  $dx = +5$ —the translation between the two images.

Fig. 4 shows the results of this experiment. Note that with normalized cross-correlation, only positive correlations (enhancement) are of interest so the graph shows the rectified (minimum value is zero) cross-correlation. As can be seen from the variations in the computed similarity measures, the three information measures show similar results and have a maximum at  $dx = 5$  (the correct displacement between the images). Normalised cross-correlation, however, does not give the correct answer. This can be explained by considering what each of the similarity measures mean. The difference between these measures is that normalized cross-correlation measures the error or *disagreement* between two locations whereas the three information measures compute the supporting evidence




Image 1







Image 2



**Normalised Cross Correlation (rectified)**



**Mutual Information**






Fig. 4. Comparison of the different similarity measures under contrast enhancement. The intensity profiles show the variation in intensity in both images along the line shown in image 1; Image 1 has a single feature to simulate a enhancing tissue type. Image 2 is a translation of image 1 by 5 pixels with a rotation of  $10^\circ$ , an enhancement of 40% and the intensity gradient reversed. The feature has 120% enhancement. The four graphs show the variation of the similarity measures between a kernel in image 2 at  $x = 20$  (indicated on image 2 by the black cross) and a kernel in image 1 moved over a displacement  $-20 < dx < 20$ .

or *agreement*. Contrast enhancement causes local areas of *disagreement* (changes in the image) where the contrast is taken up strongly, i.e., the area is potentially suspicious.

A registration algorithm based on normalized cross-correlation would attempt to reduce these areas of disagreement, as has happened in the above example, potentially causing a tumour to be missed. Consequently, mutual information is a more suitable similarity measure for breast MR image registration. This effect is confirmed by a result that will be presented in Section 9.2 where using normalized cross-correlation in the registration algorithm distorts the shape of a tumour. The choice of which of the three mutual information measures to use is almost arbitrary. We have found that in practice the differences in results are negligible.

## 6. The non-rigid registration algorithm

The previous section introduced several measures of similarity between image patches. In this section an algorithm for finding the non-rigid motion field is introduced. The algorithm uses a similarity measure expressed between two patches at  $(x, y)^T$  in image 1 and  $(x + u, y + v)^T$  in image 2 as  $S(\mathbf{r}, \mathbf{u})$  where  $\mathbf{r} = (x, y)^T$  and  $\mathbf{u} = (u, v)^T$ .

As with the majority of image registration techniques, the objective of the algorithm is to combine local information to form a globally optimal solution, while coping with the aperture problem and filtering noise.

Given image intensities in two images,  $I_0(\mathbf{r})$  and  $I_n(\mathbf{r} + \mathbf{u})$ , we model the contrast agent as a function  $F$ :

$$I_n(\mathbf{r} + \mathbf{u}) \approx F(I_0(\mathbf{r})),$$

where  $\mathbf{u}$  is the displacement between image 0 and image  $n$ .

Mutual information is a measure of the likelihood that there exists a functional relationship between  $I_0(\mathbf{r})$  and  $I_n(\mathbf{r} + \mathbf{u})$ :

$$L(I_n(\mathbf{r} + \mathbf{u}) \approx F(I_0(\mathbf{r}))) = MI(I_0(\mathbf{r}), I_n(\mathbf{r} + \mathbf{u})).$$

Consequently, the initial similarity scores, at a given location in the image, are measures of how likely it is that the second image is functionally related at that location for that motion. They can therefore be considered as estimates of the prior probabilities of those motions. Bayes' theorem can be used locally to update the estimated probability of each motion given the current motion field and a model of the deformation. Bayes' theorem can be expressed as:

$$P(Y|\mathbf{x}) = \frac{P(\mathbf{x}|Y)P(Y)}{P(\mathbf{x})},$$

where  $P(Y|\mathbf{x})$  is the probability that the answer or 'state' is  $Y$  given the measured values  $\mathbf{x}$ . On the right hand side of the equation,  $P(\mathbf{x}|Y)$  is the probability of the measurements  $\mathbf{x}$  given that the 'state' is  $Y$  (the class conditional probability).

Bayes' theorem can hence be used to compute the probability of a motion  $\mathbf{u}$  at  $\mathbf{x}$ , given the surrounding motion:

$$P(\mathbf{u}_x | \mathbf{u}_{S \setminus x}) \propto P(\mathbf{u}_{S \setminus x} | \mathbf{u}_x) P(\mathbf{u}_x), \quad (3)$$

where  $P(\mathbf{u}_x)$  is the prior probability of displacement (motion)  $\mathbf{u}$  at location  $\mathbf{x}$  and  $\mathbf{u}_{S \setminus x}$  denotes the motion flows of the surrounding locations. The denominator of the right hand side of Bayes' Theorem is a normalization factor, which results in the sum of the posterior probabilities being 1.0. The same thing can be achieved by normalizing the estimates of  $P(\mathbf{u}_x | \mathbf{u}_{S \setminus x})$  after each iteration by:

$$P_{\text{norm}}(\mathbf{u}_x | \mathbf{u}_{S \setminus x}) = \frac{P(\mathbf{u}_x | \mathbf{u}_{S \setminus x})}{\sum_x P(\mathbf{u}_x | \mathbf{u}_{S \setminus x})}. \quad (4)$$

A model of  $P(\mathbf{u}_{S \setminus x} | \mathbf{u}_x)$  is required which expresses the model of deformation. Considering only models whose conditional dependencies are *locally* dependent:

$$P(\mathbf{u}_{S \setminus x} | \mathbf{u}_x) \equiv P(\mathbf{u}_{\delta x} | \mathbf{u}_x),$$

where  $\mathbf{u}_{\delta x}$  denotes the motion flows at locations immediately adjacent to  $\mathbf{x}$ . We assume that the probability can be written as a product of potentials, which we take to be Gaussian:

$$P(\mathbf{u}_{\delta x} | \mathbf{u}_x) = \prod_{\delta x} P(\mathbf{u}_{x+\delta x} | \mathbf{u}_x).$$

At the points of interest in the image (image features) which drive the algorithm, the probability distributions of motions at a single location are sharply peaked. Consequently, for efficiency, we use the maximum as an estimate of  $P(\mathbf{u}_{x+\delta x} | \mathbf{u}_x)$  which leads to a suitable model for  $P(\mathbf{u}_{S \setminus x} | \mathbf{u}_x)$ :

$$P(\mathbf{u}_{\delta x} | \mathbf{u}_x) = \prod_{\delta x} \max_{v \in \mathcal{N}_u} \left( P(\mathbf{v}_{x+\delta x}) \exp \left( -\frac{(v-u)^2}{2\sigma_u^2} \right) \right), \quad (5)$$

where  $P(\mathbf{v}_{x+\delta x})$  is the probability of motion  $\mathbf{v}$  at  $\mathbf{x} + \delta \mathbf{x}$  and  $\mathcal{N}_u$  is the set of motions similar to  $\mathbf{u}$  (the set of motions such that  $|v - u| < \varepsilon$ ). The model in Eq. (5) considers each neighbouring location  $\mathbf{x} + \delta \mathbf{x}$ . For each neighbouring location, the maximum probability of similar motions  $\mathbf{u} + \delta \mathbf{u}$ , weighted by a Gaussian, is taken. The probability of the surrounding motions given  $\mathbf{u}$  as the correct motion is estimated as the product of these maxima from each location. A Gaussian distribution is a reasonable model to use here, given the stated *optimum* solution is that which minimizes  $\nabla^2 \mathbf{u}$  and  $\mathbf{u}$  being a Gaussian is a solution to the diffusion equation.

A brief note on implementation: The update algorithm requires multiplications and taking maxima. It is therefore computationally cheaper to take logarithms of the similarity data so that all multiplication operators become additions. The model equations then become:

$$P_{\log}(\mathbf{u}_{\delta x} | \mathbf{u}_x) = \sum_{\delta x} \max_{v \in \mathcal{N}_u} \left( P_{\log}(\mathbf{v}_{x+\delta x}) - \left( \frac{(v-u)^2}{2\sigma_u^2} \right) \right) \quad (6)$$

and the update equation becomes:

$$P_{\log}(\mathbf{u}_x | \mathbf{u}_{S \setminus x}) \propto P_{\log}(\mathbf{u}_{S \setminus x} | \mathbf{u}_x) + P_{\log}(\mathbf{u}_x). \quad (7)$$

After applying Eq. (7) iteratively, the resulting motion field can be computed by taking, as the estimate of the motion at each sampled point in the image, the motion corresponding to the maximum posterior probability:

Table 1  
Bayesian registration algorithm

- 
- (1) Compute similarity scores at a set of equally spaced sample locations across the image for all possible motions according to Eq. (2).
  - (2) Normalize the set of similarities at each location to estimate the prior probabilities  $P(\mathbf{u}_x)$  using Eq. (4).
  - (3) Take natural logarithms of the prior probabilities.
  - (4) Compute the posterior probability of each motion given its surrounding estimates using the update Eq. (7).
  - (5) Re-normalize the estimated probabilities. Goto step (4) unless the required number of iterations have been completed. The number of iterations is dependent on the distance information must propagate to interpolate between features (see Section 6.1 for discussion).
  - (6) Find the motion field by taking, as the estimate of the optic flow at each location, the displacement corresponding to the maximum posterior probability using Eq. (8).
- 

$$\mathbf{u}(x) = \max_{P(\mathbf{u}_x | \mathbf{u}_{S \setminus x})} \{\mathbf{u}_x\}. \quad (8)$$

This method is very similar to Markov random field methods. The difference is that normally, Markov random field algorithms estimate a set of discrete labels. Consequently a model is required (often taken to be a Gaussian) of the prior distribution given a particular estimated label. At each iteration,  $P(\mathbf{u}_x | \mathbf{u}_{S \setminus x})$  is computed for all  $\mathbf{u}_x$  and the new estimated label (or in our case motion) is the one with the greatest probability.

The complete algorithm is summarized in Table 1.

### 6.1. Discussion

The types of deformations allowed by the algorithm are implicit in the model for  $P(\mathbf{u}_{\delta x} | \mathbf{u}_x)$ . The model is designed to minimize  $\delta u / \delta x$ . The behaviour of the algorithm is influenced by the following parameters:

- The variance  $\sigma_u$  in  $P(\mathbf{u}_{\delta x} | \mathbf{u}_x)$  controls the probability of  $\delta u$  in  $\delta u / \delta x$ .
- The grid sampling frequency (the spacing over which motion estimates are taken) scales  $\delta x$  in  $\delta u / \delta x$ . Hence, for a given sampling frequency  $\delta x$ , the variance  $\sigma_u$  can be computed to control  $\delta u / \delta x$ .

The similarity scores are computed at a set of equally-spaced locations across the images. The number of iterations required to compute the posterior probabilities depends on the sample frequency of the motion flow estimates.

### 6.2. Computing a parametric transformation from the motion flow field

A motion flow field between a pre and post-contrast image pair is a set of motion estimates at discrete locations in the image. In order to correct for movement, the motion at every pixel must be evaluated. A simple method would be to compute the motion at a pixel using bi-linear interpolation of the nearest surrounding motion vectors. A better method is to fit a parametric transformation to the motion flow field which is then continuous and can be evaluated at any location.



Declerck's approximating B-spline mesh algorithm [5] fits a cubic tensor product B-spline mesh of the form:

$$u(x, y) = \sum_{i=0}^{n_x-1} \sum_{j=0}^{n_y-1} \alpha_{i,j} B_{i,K}^x(x) B_{j,K}^y(y), \quad (9)$$

$$v(x, y) = \sum_{i=0}^{n_x-1} \sum_{j=0}^{n_y-1} \beta_{i,j} B_{i,K}^x(x) B_{j,K}^y(y) \quad (10)$$

to a set of displacements subject to an energy minimization with accuracy and smoothing terms.  $n_x$  is the number of control points in the  $x$ -direction and  $B_{[i,j],K}^{[x,y]}$  are the cubic B-spline basis functions.

The fitting is by minimization of an energy function:

$$J^x(u) = J_{\text{position}}^x(u) + J_{\text{smooth}}^x(u)$$

and analogously for  $J^y(u)$  where

$$J_{\text{position}}^x(u) = \sum_{i=1}^N (u(x_i^1, y_i^1) - x_i^2)^2, \quad (11)$$

$$J_{\text{smooth}}^x(u) = \rho_{\text{smooth}} \iint_{\mathbb{R}^2} [u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2]. \quad (12)$$

The integrals required for Eq. (12) are calculated with the Gauss–Legendre algorithm and the minimization is performed with respect to each of the parameters  $\alpha_{i,j}$  and  $\beta_{i,j}$  using conjugate-gradient descent. A full derivation and explanation of this method is given by Declerck in [5] and [6].

## 7. Multi-scale implementation

The algorithm described in the previous section for computing the posterior probabilities is of order  $O(n^2)$  complexity with the scale of the search window (i.e., doubling the maximum movement allowable in the algorithm, quadruples the time taken) and also of order  $O(n^2)$  complexity with respect to the motion estimates sample frequency. To alleviate this problem, the algorithm was implemented as a coarse-to-fine multi-scale algorithm. The initial iteration computes the approximate transformation, sub-sampling the images and sub-sampling the possible displacements. Subsequent iterations use progressively finer sampling of the images and displacements.

The complete multi-scale algorithm is summarized in Table 2.

Table 2

## Multi-scale registration algorithm

- 
- (1) Compute similarity scores at a set of locations across the image for sampled possible displacements according to Eq. (2).
  - (2) Normalize the set of similarities at each location to estimate the prior probabilities  $P(u_x)$  using Eq. (4).
  - (3) Take natural logarithms of the prior probabilities.
  - (4) Compute the posterior probability of each optic flow given its surrounding estimates using the update Eq. (7).
  - (5) Re-normalize the estimated probabilities. Goto step (4) unless the required number of iterations have been completed.
  - (6) Find the motion field by taking, as the estimate of the optic flow at each location, the displacement corresponding to the maximum posterior probability using Eq. (8) and compute a parametric transformation between the images (Section 6.2).
  - (7) Transform the image  $I_2^n$  to obtain the motion corrected image  $I_2^{n+1}$ .
  - (8) Goto step (1) and repeat at the next finer scale.
- 

## 8. Compensating for the bias field

In certain cases, MR image intensities,  $I(x, y, z)$ , are corrupted by a substantial ‘bias’ field,  $B(x, y, z)$  [11,34].

$$I(x, y, z) = f_{\text{MR}}(x, y, z)B(x, y, z) + n(x, y, z),$$

where  $n(x, y, z)$  is additive noise. Fig. 5 illustrates the variation in MR signal level across of region of fat in the breast. We have developed a model-based technique to correct for the bias field which is described in [12].

When the breast moves, its position in the bias field  $B(x, y)$  changes. Consequently, even without enhancement, there is an expected intensity change. In the presence of movement, the apparent contrast enhancement is computed as:

$$\text{CE}_{\text{apparent}} = \frac{I_{\text{post}}(x + u, y + v)}{I_{\text{pre}}(x, y)}.$$

The true enhancement is

$$\text{CE}_{\text{true}} = \frac{I_{\text{post}}(x + u, y + v)}{B(x + u, y + v)} \bigg/ \frac{I_{\text{pre}}(x, y)}{B(x, y)} = \frac{I_{\text{post}}(x + u, y + v)}{I_{\text{pre}}(x, y)} \cdot \frac{B(x, y)}{B(x + u, y + v)}.$$

The effective change in post-contrast signal level due to movement through the bias field (the signal change in  $I_{\text{post}}$  with respect to the difference between the apparent and true contrast enhancement) is therefore:

$$\Delta I = I_{\text{post}}(x + u, y + v) - I_{\text{post}}(x + u, y + v) \frac{B(x, y)}{B(x + u, y + v)} \quad (13)$$

$$= I_{\text{post}}(x + u, y + v) \left( 1 - \frac{B(x, y)}{B(x + u, y + v)} \right). \quad (14)$$

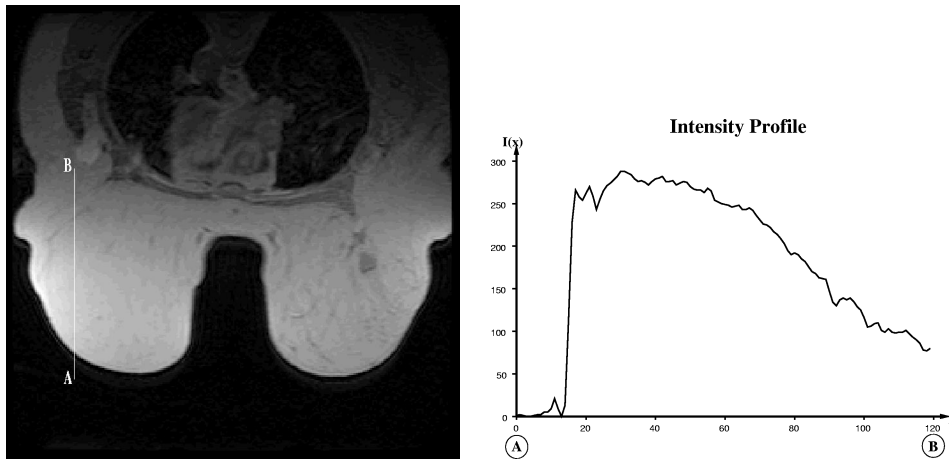



Fig. 5. A typical MRI image. The effect of the bias field is clearly apparent as a signal decrease towards the chest wall. The breast is almost entirely fat, so the breast tissue is constant and the MR signal intensity should ideally also be constant; the intensity variation along the line  $AB$ , however, shows a 60% signal attenuation through the breast.

For registration, the intensity change due to movement through the bias field can be considered as an enhancement (although this may be negative). It is important to note this effect, however, since the purpose of motion correction is to compute the enhancement in the breast and the effect of the bias field on the motion-corrected images will result in an apparent increase or decrease in the amount of contrast enhancement. Consequently, whether or not the images are bias-corrected before *computing* the motion field (it makes very little difference to the performance of the algorithm) they must be bias-corrected before *correcting* for the computed motion field, or at least compensated for the motion through the bias field<sup>2</sup> using Eq. (14).


## 9. Testing

This section will present results on three patients that moved during the MR scan and show that, using the registration algorithm described in this chapter, the motion can be corrected for. One further result is given to confirm the argument for the use of mutual information as the similarity measure rather than normalized cross-correlation. We finish by presenting a method for quantifying the performance of the algorithm on real images and apply this to the third example which exhibited the largest motion.


<sup>2</sup> Bias correcting the images involves correcting each image in the sequence for the bias field. It is only the enhancement that is of diagnostic importance so as an alternative, the post-contrast images can be corrected for the change in position in the bias field due to motion.




(a)




(b)



(c)



(d)



(e)



(f)




Fig. 6. Results of the registration algorithm applied to pre- and post-contrast MR images from a patient who moved during the scan. (a) Pre-contrast image. (b) Post-contrast image. (c) Subtraction image. (d) Motion field. (e) Corrected post-contrast image. (f) Corrected subtraction image. Images (a), (b) and (c) show the pre-, post-contrast and subtraction images, respectively. Image (d) shows the computed motion field. Images (e) and (f) show the corrected image and the corrected subtraction image. The intensity profiles show the variation in intensity along the line shown in image (a). These demonstrate the accuracy of the registration at the left breast edge. The horizontal displacement between the resulting breast edges is less than half a pixel.

### 9.1. Results

The patient in Fig. 6 moved by approximately 10 mm. The image sequence was such that NRM (author) reported that the scan sequence was '*near impossible to interpret*'. There was nothing clinically significant in the scan, though this cannot be determined from the subtraction images or enhancement graphs, and was difficult to determine by visual observation alone. The intensity profiles in Fig. 6 verify that there is enhancement on the left hand side of the breast which can be seen in the corrected subtraction image and show that the breast edge has been registered to an accuracy of less than half a pixel.

In Fig. 7, the patient has a tumour which is circled in the subtraction image. The patient has moved approximately 2 mm between pre and post-contrast images. The surface plots show the intensity surfaces of the subtraction images in the region around the tumour with and without motion correction. The peak in the surface corresponds to the tumour and is more strongly peaked in the corrected subtraction surface plot. This implies that even though the motion is very slight, after applying the registration algorithm, the tumour is better localized in the motion corrected image.

In Fig. 8, the registration algorithm has been used to correct the post-contrast image for a patient who relaxed her pectoral muscles between the pre and post-contrast acquisitions. The movement is extremely non-rigid (the pectoral muscle<sup>3</sup> changes shape and moves approximately 15 mm). The patient has a tumour which also moves 7 or 8 mm as a result of the movement within the breast. It is interesting to note that in this case, the movement is almost confined within the breast, with the breast edge moving only about 3 mm. The registration algorithm correctly identifies the motion field and the tumour in the corrected post-contrast image has been moved back to the correct position. The largest error in the corrected post-contrast image is at the right hand edge of the pectoral muscle where there is a 1-pixel error between the edge in the pre and corrected post-contrast image. The split images are formed by displaying the pre-contrast image in the top left and bottom right quadrants<sup>4</sup> of the image and the post-contrast image in the bottom left and top right quadrants. These images show motion errors at the quadrant boundaries. In the motion corrected split image, the boundaries of the pectoral muscle and fibroglandular tissue are aligned demonstrating that the algorithm has correctly registered the images.

### 9.2. Comparison of similarity measures

In Section 5 the choice of mutual information over normalized cross-correlation as the similarity measure was made on the basis of mutual information measuring the amount of agreement, whereas cross-correlation measures the amount of disagreement. It was noted that the choice makes very little difference to the performance of the algorithm in the majority of cases. The results in Fig. 9, however, show one example where there is a marked difference between the similarity measures.

Comparing the results for normalized cross-correlation and mutual information in Fig. 9, it can be seen that with normalized cross-correlation, the algorithm has computed the

---

<sup>3</sup> The pectoral can be identified as the lower intensity tissue between the breast and the chest wall.

<sup>4</sup> The quadrants are delineated by the cross.







Fig. 7. Results of the registration algorithm applied to pre- and post-contrast MR images from a patient who moved during the scan. Images (a) and (b) show the pre-contrast and subtraction images, respectively. There is a tumour present which is circled in the subtraction image. Image (c) shows the corrected subtraction image. The surface plots show the intensity surfaces from the subtraction images within the region of interest (circled in image (b)).




(a)




(b)




(c)



(d)



(e)



(f)







Fig. 8. Results of the registration algorithm applied to pre- and post-contrast MR images from a patient who moved during the scan. (a) Pre-contrast image. (b) Post-contrast image. (c) Subtraction image. (d) Motion field. (e) Corrected post-contrast image. (f) Corrected subtraction image. (g) Pre/post split image. (h) Corrected pre/post split image. Images (a), (b) and (c) show the pre-, post-contrast and subtraction images, respectively. Image (d) shows the computed motion field. Images (e) and (f) show the corrected image and the corrected subtraction image. Images (g) and (h) are split between pre- and post-contrast images. The pre-contrast image is to the top left and bottom right quadrants, the post-contrast image in the bottom left and top right quadrants. Image (g) is split between the original pre- and post-contrast images; image (h) between the pre- and corrected post-contrast images.



Normalized cross-correlation

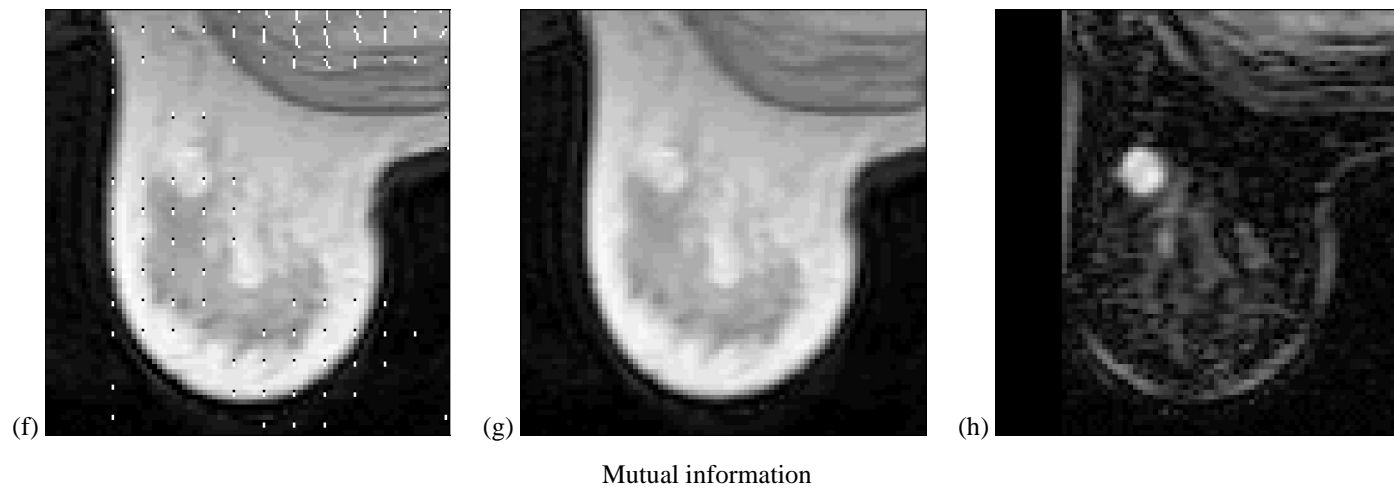


Fig. 9. A comparison of the registration fields computed with different similarity measures. (a) Pre-contrast image. (b) Post-contrast image. (c) Motion field. (d) Corrected post-contrast image. (e) Corrected subtraction image. (f) Motion field. (g) Corrected post-contrast image. (h) Corrected subtraction image. Images (a), (b) show the pre- and post-contrast images. Images (c), (d) and (e) show the computed motion field, corrected and corrected subtraction images with normalized cross-correlation. Images (f), (g) and (h) show the same with mutual information.

motion incorrectly in the region around the tumour, precisely the location which it is required to be most accurate. Using mutual information, no such problems occur which confirms the argument for its use as the similarity measure in contrast-enhanced breast MR images.

### 9.3. Consistency of solution

Although the algorithm appears to produce visually good results, in order to assess the performance of the algorithm a quantitative measure of its accuracy is required. It is difficult to establish the *true* transformation between a pre and post-contrast image for a real MR scan and artificial images or phantoms will produce artificially good results. No realistic breast phantoms are currently available; although it would perhaps be possible to use a cadaver as a phantom, into which markers could be implanted to measure motion, this is not seen as ethically feasible.

An alternative measure of the algorithm's performance, is the consistency of the forward and backward transformations computed by the algorithm.

$$T_1 = \text{Register}(\text{Image}_1, \text{Image}_2),$$

$$T_2 = \text{Register}(\text{Image}_2, \text{Image}_1),$$

so  $T_1$  is the transformation that maps Image 2 onto Image 1 and  $T_2$  is the transformation that maps Image 1 onto Image 2. If the registration algorithm is accurate then:

$$T_1 = T_2^{-1}$$

and a measure of the algorithm's consistency over the breast area,  $A$ , is

$$C = \frac{1}{A} \int_A (T_2(T_1(x)) - x) dA.$$

To investigate the consistency measure, it was applied to a pair of artificial test images. The correct transformation between the images is a 20° rotation about the centre of the image. The consistency error can therefore be compared across the image to the true registration error and the results of this comparison are shown in Fig. 10. Note that, although artificial images cannot be used to assess the performance of the registration algorithm itself, they can be used to investigate the consistency measure, where it is the difference between the true errors and consistency errors that is of interest, not the errors themselves.

From Fig. 10 the consistency errors can be compared with the true errors. Two conclusions can be drawn from this experiment. Firstly, the consistency errors are, on average, greater than the true errors and secondly, though the errors are not well correlated, it is likely that an area which shows low consistency errors will be accurately registered and vice versa.

In order to examine the performance of the algorithm on real images, the consistency measure was applied to the images of Fig. 8 which exhibited the most complex motion. The results are shown in Fig. 11. As is to be expected the worst errors are around the right hand edge of the pectoral muscle which moved approximately 15 mm. Note however, that

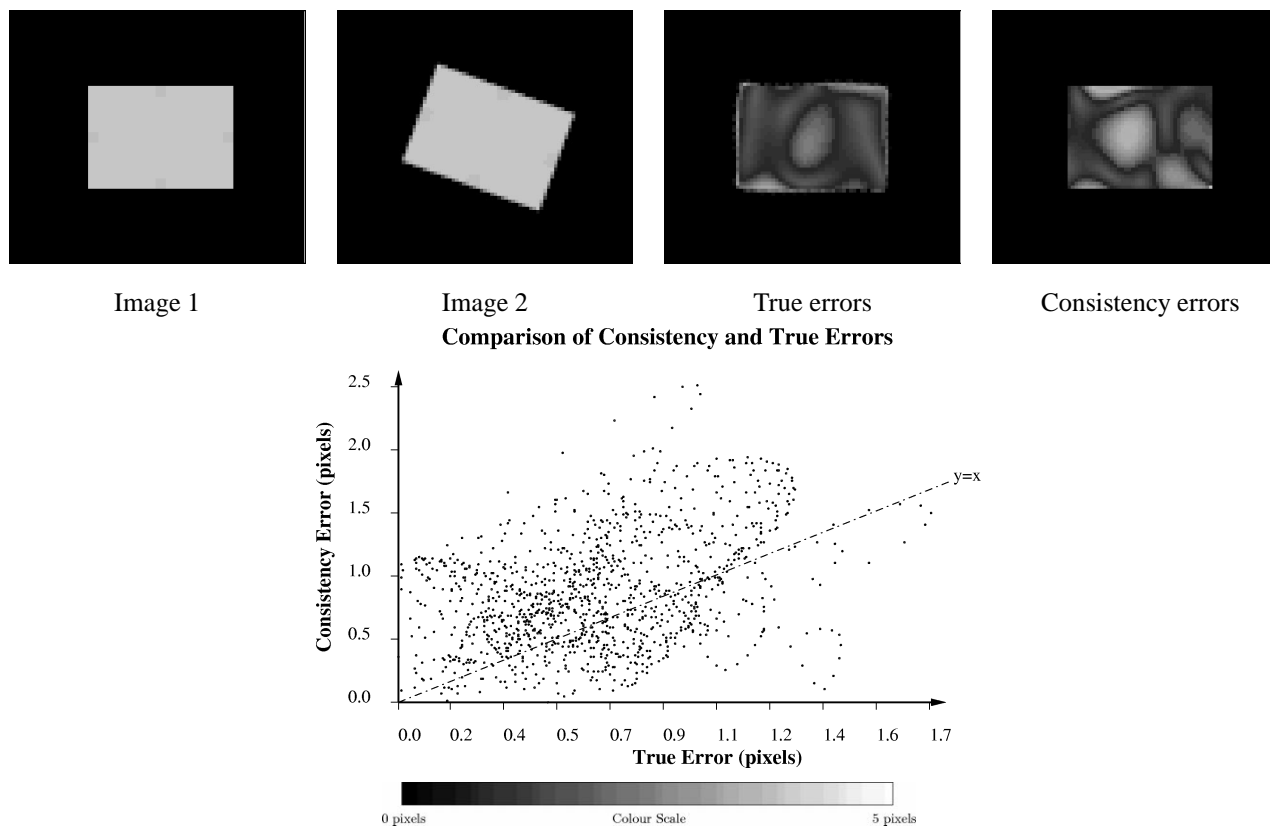


Fig. 10. Applying the consistency measure to a pair of artificial images. The colour plots show the spatial variation of true and consistency errors across the images; the largest registration errors are in the centre of the image (the centre of rotation) which is also the location of the largest consistency errors. The scatter plot shows how the consistency errors relate to the true errors.

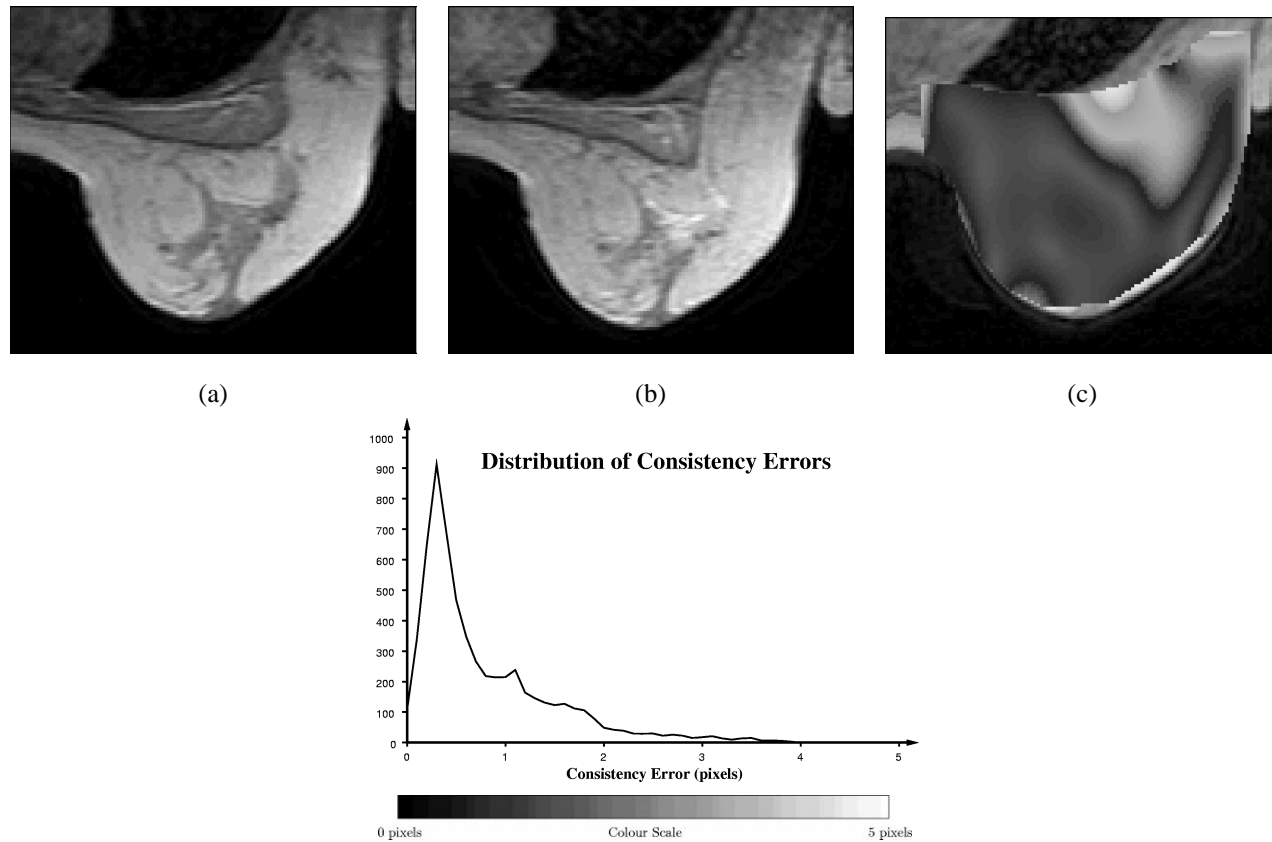


Fig. 11. Applying the consistency measure to the images from Fig. 8. (a) Pre-contrast image. (b) Post-contrast image. (c) Registration consistency. Images (a) and (b) show the pre- and post-contrast images, respectively. Image (c) shows the spatial variation of registration consistency over the breast. The graphs shows the distribution of consistency values.

although the colour scale shows the area of inconsistency, the magnitude is only about 2 pixels which indicates a considerable improvement due to the registration algorithm. There is a small area at the top of the pectoral muscle where the inconsistency is 4 pixels. This region appears to have a different texture in the two images, possibly due to noise, and is of little clinical significance.

## 10. Conclusion

This paper has addressed one of the key problems in dynamic breast MR imaging, that of image registration. An algorithm capable of robustly correcting for patient movement has been presented and its performance demonstrated on several sets of images. The algorithm avoids the use of image derivatives which tend to be severely affected by noise, but uses a novel algorithm for computing the global motion field using all the information.

The algorithm described in this chapter can easily be extended for use on 3D data sets. There were not, however, any 3D dynamic breast sequences available for this project. An alternative would be to link the motion fields between neighbouring slices. This could be achieved by changing the model of deformation to apply across slices and computing the motion fields for all slices in parallel.

The model of deformation used in the registration algorithm was a simple one. The Laplacian model is a simplification of the linear elastic model, which in itself would be a simplification of reality. Estimates have been made of the Young's Modulus of the different tissue types in the breast [24] and these could be used to develop a more realistic deformation model though this would also require accurate segmentation of tissue types.

## References

- [1] J.L. Barron, D.J. Fleet, S.S. Beauchemin, Performance of optical flow techniques, *Internat. J. Computer Vision* 12 (1) (1994) 43–77.
- [2] M. Bro-Nielsen, Medical image registration and surgery simulation, Ph.D. Thesis, Department of Mathematical Modelling, Technical University of Denmark, 1996.
- [3] G.E. Christensen, R.D. Rabbitt, M.I. Miller, 3D brain mapping using a deformable neuroanatomy, *Physics in Medicine and Biology* 39 (1994) 609–618.
- [4] S. Ciatto, L. Cataliotti, V. Distanti, Nonpalpable lesions detected by mammography: Review of 512 consecutive cases, *Radiology* 165 (1987) 99–102.
- [5] J. Declerck, G. Subsol, J.P. Thirion, N. Ayache, Automatic retrieval of anatomical structures in 3D medical images, Technical Report, INRIA, 1995.
- [6] J. Declerck, G. Subsol, J.P. Thirion, N. Ayache, Automatic retrieval of anatomical structures in 3d medical images, in: N. Ayache (Ed.), *First International Conference on Computer Vision, Virtual Reality and Robotics in Medicine, CVRMed-95, Nice, France, Lecture Notes in Computer Science*, Springer, Berlin, 1995.
- [7] F. Devernay, Vision stéréoscopique et propriétés différentielles des surfaces, Ph.D. Thesis, L'École Polytechnique, 1997.
- [8] J. Esteve, A. Krickler, J. Ferlay, D.M. Parkin, Facts and figures of cancer in the European community, Technical Report, International Agency for Research on Cancer, Lyon, France, 1993.
- [9] S. Feig, Breast masses: Mammographic and sonographic evaluation, *Radiol. Clin. North Amer.* 30 (1992) 67–92.

- [10] S. Gong, M. Brady, Parallel computation of optic flow, in: O. Faugeras (Ed.), European Conference on Computer Vision I, Springer, Berlin, 1990, pp. 124–133.
- [11] R. Guillemaud, M. Brady, Enhancement of mr images, in: K.H. Höhne, R. Kikinis (Eds.), Visualisation in Biomedical Computation (VBC-96), Lecture Notes in Computer Science, Vol. 1131, Springer, Berlin, 1996, pp. 107–116.
- [12] P.M. Hayton, Analysis of contrast-enhanced breast MR images, Ph.D. Thesis, Oxford University, Department of Engineering, Oxford, 1998.
- [13] S.H. Heywang-Koebrunner, Nonmammographic breast imaging techniques, *Current Opinions in Radiology* 4 (1992) 146–154.
- [14] R.P. Highnam, M. Brady, Mammographic Image Processing, Kluwer Series in Computer Vision, Kluwer, Dordrecht, 1998.
- [15] E.C. Hildreth, Computations underlying the measurement of visual motion, *Artificial Intelligence* 23 (1984) 309–354.
- [16] B.K.P. Horn, B.G. Schunck, Determining optical flow, *Artificial Intelligence* 17 (1981) 185–203.
- [17] D. Ikeda, I. Anderson, Ductal carcinoma in situ: Atypical mammographic appearances, *Radiology* 172 (1989) 661–666.
- [18] V. Jackson, K. Dines, L. Bassett, R. Gold, H. Reynolds, Diagnostic importance of the radiographic density of noncalcified breast masses: Analysis of 91 lesions, *Amer. J. Rentgenology* 157 (1991) 25–28.
- [19] J.B.A. Maintz, M.A. Viergever, A survey of medical image registration, *Medical Image Analysis* 2 (1) (1998) 1–36.
- [20] H.-H. Nagel, On the estimation of optic flow: Relations between different approaches and some new results, *Artificial Intelligence* 33 (1987) 299–324.
- [21] H.-H. Nagel, Extending the ‘oriented smoothness constraint’ into the temporal domain and the estimation of derivatives of optical flow, in: Proc. European Conference on Computer Vision I, 1990, pp. 139–148.
- [22] B. Porter, V. Taylor, J. Smith, V. Tsao, Contrast-enhanced magnetic resonance mammography, *Acad. Radiology* 1 (3) (1994) S36–S46.
- [23] P. Pretorius, P. Hayton, N. Moore, M. Brady, A clinical assessment of the use of a model of contrast enhancement for analysis of dynamic MR breast images, in: Proc. Medical Image Understanding and Analysis ’97, 1997, pp. 57–60.
- [24] P. Pretorius, N. Moore, P. Hayton, M. Brady, The use of compression during MR imaging of the breast to demonstrate elasticity variation between different tissues and lesions, in: Imaging: Abstracts from Symposium Mammographicum ’98, Vol. 10(1), York, Science Technology and Letters, 1998.
- [25] E. Sickles, Mammographic features of 300 consecutive nonpalpable breast cancers, *Amer. J. Rentgenology* 146 (1986) 661–663.
- [26] E. Simoncelli, E. Adelson, D. Heeger, Probability distributions of optical flow, in: CVPR, Hawaii, Inst. Elec. Eng., 1991, pp. 310–315.
- [27] S.M. Smith, ASSET-2: Real-time motion segmentation and object tracking, Internal Technical Report TR95SMS2, Defence Research Agency, Chobham Lane, Chertsey, Surrey, UK, 1995. Available at [www.fmrib.ox.ac.uk/~steve](http://www.fmrib.ox.ac.uk/~steve) for downloading.
- [28] S.M. Smith, J.M. Brady, ASSET-2: Real-time motion segmentation and shape tracking, *IEEE Trans. Pattern Anal. Mach. Intell.* 17 (8) (1995) 814–820.
- [29] C. Studholme, Measures of 3D medical image alignment, Ph.D. Thesis, University of London, 1997.
- [30] J. Thirion, Fast non-rigid matching of 3D medical images, Technical Report 2547, INRIA, 1995.
- [31] B.C. Vemuri, H. Shuangying, D. Sahni, C. Leonard, C. Mohr, R. Gilmore, J. Fitzsimmons, An efficient motion estimator with application to medical image registration, *Medical Image Analysis* 2 (1) (1998) 79–98.
- [32] P. Viola, W. Wells, Alignment by maximization of mutual information, in: IEEE Computer Society (Ed.), ICCV, 1995, pp. 16–23.
- [33] W.M. Wells III, W.E.L. Grimson, R. Kikinis, F.A. Jolesz, Adaptive segmentation of MRI data, *IEEE Trans. Medical Imaging* 15 (4) (1996) 429–442.
- [34] W.M. Wells III, P. Viola, H. Atsumi, S. Nakajima, R. Kikinis, Multi-model volume registration by maximization of mutual information, *Medical Image Analysis* 1 (1) (1996) 35–51.