

Non-rigid Multimodal Image Registration Using Local Phase

Matthew Mellor^{*} and Michael Brady

Medical Vision Laboratory, University of Oxford, UK
{matt,jmb}@robots.ox.ac.uk

Abstract. Non-rigid registration of multimodal images is a challenging problem. One approach, maximization of mutual information, has been shown to be effective for registering certain image modalities and is currently considered the standard against which all other techniques are measured. In this paper, we propose an alternative representation of an image based on local phases rather than intensities; we then show how mutual information can be extended to this representation. Local phase acts as a description of local image structure, enabling mutual phase information to detect complex image relationships that are hard or impossible to detect using mutual intensity information. Typical results are presented, comparing the performance of phase and intensity mutual information methods on simulated MR and ultrasound images.

1 Introduction

This paper illustrates the effectiveness of a representation of an image, based on local phase, which significantly enhances non-rigid registration of a wide variety of multimodal images. Since it has become the reference against which other measures and non-rigid registration algorithms are judged, the experiments in this paper use mutual information (MI) of local phases to maximise a *structural* relationship between images. It should be understood from the outset, however, that the statistical behaviour of phase is highly predictable and so it is theoretically possible to design multimodal non-rigid registration algorithms which employ substantially more prior information than is available to conventional intensity-based mutual information methods. Several methods for non-rigid registration of images based on mutual information of intensities have been introduced [1,2,3]. These methods use the global intensity statistics to inform local alterations to the registration. Regardless of the specific method used, these methods are only able to register images for which the global intensity statistics are an accurate reflection of the local statistics. Unfortunately, but significantly, this is manifestly not the case for many modalities, potentially including some of those for which rigid registration using MI has been asserted to be feasible.

^{*} Matthew Mellor is funded by the EPSRC as part of the MIAS-IRC. (GR/N14248)

2 Local Phase and Energy

Local phase is a qualitative description of the shape of an image location. It may be thought of as an amplitude-weighted average phase of the (windowed) Fourier components at each point in the signal. At a locally anti-symmetric point in the signal, the Fourier components are also predominantly anti-symmetric (i.e. sine functions), whereas at locally symmetric points they are predominantly symmetric (cosine functions). The local phase is often considered in conjunction with a second quantity, called the local energy. Local energy can be thought of as the combined energy, after interference, of the Fourier components and acts as a measure of local signal activity. Local phase has a number of important properties which make it a natural choice for image registration. First, the local variation of the phase is well approximated by the phase gradient, which has been shown to increase the accuracy of matching algorithms [4]. Second, the marginal distributions, the treatment of which distinguishes the various information theoretic methods (mutual information, joint entropy, normalised mutual information etc) are highly predictable for phase. Figure 1 shows typical intensity and phase distributions for an MR-T1 image. Whereas the phase statistics are relatively flat, the intensity statistics clearly show the presence of several different classes. In general, an intensity only has meaning when interpreted in the context of other pixels in the image. In contrast, phase is a fundamental property, requiring no context for interpretation and this is reflected in the marginal. From the perspective of mutual information, it may be argued that a phase marginal contains limited information, relative to a typical intensity marginal.




Fig. 1. Typical marginal distributions of an MR-T1 image: intensity (solid) is complex, while phase (dotted) is relatively flat

There is more than one way to extend the notion of phase to two dimensions. One possibility is to estimate the phase of the one dimensional signal in the direction of maximal local energy using steerable filters. For the results presented

here, the local phase and energy have been estimated using the monogenic signal, described below. This choice was made in part for reasons of computational speed and simplicity, in part for mathematical elegance. However, in principle, there is no reason why the results presented here are particular to a method of phase estimation.

The monogenic signal [5] is a multi-dimensional extension to the analytic signal, which consists, at each image location, of three quantities: local energy, local phase, and local orientation. Space does not permit a detailed discussion of local energy and orientation, which can be found in [6]

The monogenic signal may be computed from the output of three filters. First, a rotationally symmetric, zero-mean filter is applied to the image, to give a bandpass image I_b : this constitutes the even component of the signal. The odd component is composed of the response of two anti-symmetric filters to the even part. These two filters, h_1 and h_2 , are described in the Fourier domain by:

$$H_1(u, v) = \frac{u}{\sqrt{u^2 + v^2}}, H_2(u, v) = \frac{v}{\sqrt{u^2 + v^2}} \quad (1)$$

where u, v are the Fourier domain coordinates. Local phase, ϕ and local energy, E , can then be calculated from these filter responses and the bandpass image I_b as follows:

$$\phi(x, y) = \tan^{-1} \left(\frac{I_b}{\sqrt{(h_1 \otimes I_b)^2 + (h_2 \otimes I_b)^2}} \right) \quad (2)$$

$$E(x, y) = \sqrt{I_b^2 + (h_1 \otimes I)^2 + (h_2 \otimes I_b)^2} \quad (3)$$

$$(4)$$

3 Filter Choice

The measures of phase estimated using the method described above depend on the particular filters used to derive the bandpass image I_b . Criteria for the choice of bandpass filters for 1D signals have been studied extensively. However, the behaviour of a two dimensional filter to a two dimensional signal is more complex than the one dimensional case. For example, when a filter, such as a difference of Gaussians, is applied to a curved image structure, the blurring causes the phase contours to migrate from the feature towards the inside of the curve. This causes a problem for registration, since blurring by such filters is not affine invariant; this makes it difficult to predict phase after a local affine transform. To avoid this, we have developed filters which exhibit very little blurring:

$$f(r) = \frac{A}{r^{\alpha+\beta}} - \frac{B}{r^{\alpha-\beta}} \quad (5)$$

where r is the distance from the centre of the filter, α is a design parameter and β is a second parameter, which must be very small compared to α to avoid blurring. A and B are chosen to give zero mean and an appropriate amplitude.

The value of the filter is infinite for $r = 0$. To avoid having an infinite central value, the discrete versions of the filter must be calculated with the average of r over the pixel rather than the central value. For a more detailed discussion of these filters and their use see [6].

4 Algorithm

The results presented here are based on a method derived from that of Crum [2]. The algorithm deforms one image, the source, to match the second image, the target. At each iteration, the present estimate of the registration is applied to the source image and its local phase and energy are estimated using the monogenic signal, and using five different members of the filter family described in the previous section. These five filters are defined by $\alpha = [2.75, 3.25, 3.75, 4.25, 4.75]$ and $\beta = 0.25$ in each case. The joint distribution of the source image phases and target images phases for each filter is estimated. This proceeds by histogram binning with 100 bins per phase, followed by smoothing with a Gaussian of width $\sigma = 7$ bins. Only phase values with an energy greater than 5% of the maximum are binned, to reduce the effects of noise. Although the resolution of the joint distributions is quite low, because of the relatively small number of pixels typically available in 2D images, any increase in resolution decreased robustness. We anticipate that this effect will be less severe for 3D images. Once the joint distribution has been estimated, the mutual information forces are computed according to equation 6, for each of the phase images estimated with the five bandpass filters as follows:

$$F(x) = K(\log[P(\phi_1(x+1), \phi_2(x))/P(\phi_1(x+1))] - \log[P(\phi_1(x-1), \phi_2(x))/P(\phi_1(x-1))]) \quad (6)$$

where $\phi_1(x)$ is the local phase at position x in the source image, K is the timestep, and $\phi_2(x)$ is the phase in the target image. To regularise, the displacement field is convolved with a Gaussian ($\sigma = 4$) at each iteration.

For the purposes of comparing phase and intensity methods fairly, an intensity algorithm was designed, similar in as many respects as possible to the phase algorithm. The intensity is used directly, with no phase estimation. The time-step for each algorithm has been hand optimised such that both methods converge at, as near as possible, the same rate.

5 Results

The phase and intensity algorithms described above were applied to sets of synthetically generated two dimensional images. The first set consists of one MR-T2 image and thirty MR-T1 images to which were applied randomly generated warps. The second set consists of the same thirty MR-T1 images, but with the addition of a slowly varying, randomly generated, multiplicative bias field. The third set consists of thirty simulated ultrasound images. These were generated by using the randomly warped MR-T1 images as the scatter density field in a simple ultrasound simulation algorithm.

5.1 MR-T1 to MR-T2 Registration

The synthetic images from which the image pairs are generated were taken from the Brainweb simulated brain database [7]. Thirty deformed T1 images were generated by applying random warps to the original T1 image. The x and y components of the warps were each generated by convolving an IID normally distributed displacement field with a Gaussian of width $\sigma = 12$ pixels. The resulting warps were normalised such that the r.m.s displacement was three pixels. An example warped T1 image is shown in figure 2.




Fig. 2. T1-MR image before (left) and after (right) application of a randomly generated non-rigid deformation

Since the deformations are created using the same smoothing method as the regulariser, they should be very compatible with the registration algorithms: any failure should be due to the similarity measure and/or minimisation strategy, not the inability of the regulariser to produce the true deformation.




Fig. 3. mean displacement errors in pixels after registration with intensity (left) and phase (right)

The average performance of the two algorithms on the first synthetic data set is illustrated in figure 3. Both methods perform well on the brain, achieving sub-

Table 1. Average registration errors in pixels for the phase and intensity methods

		Phase error		Intensity error	
		average	S.D.	average	S.D.
Brain & Skull	r.m.s	1.81	0.13	1.63	0.12
	mean	1.41	0.08	1.34	0.08
Brain	r.m.s	1.39	0.15	1.45	0.13
Alone	mean	1.11	0.11	1.21	0.09

Table 2. Average registration errors in pixels for both methods with bias field

		Phase error		Intensity error	
		average	S.D.	average	S.D.
Brain & Skull	r.m.s	1.86	0.13	1.76	0.13
	mean	1.46	0.08	1.45	0.09
Brain	r.m.s	1.42	0.15	1.66	0.13
Alone	mean	1.14	0.11	1.38	0.12

pixel accuracy for a significant proportion, with the phase method performing better on average. The likely reason for this is the greater accuracy of linear interpolation for phase values [4]. In the region of the skull however, the intensity method outperforms phase. Since structures that are parallel in both images will create a mode in the phase joint histogram, it is difficult for phase based method to register them correctly, and this is the likely reason for the error in the skull region. The exact values for the average registration quality are given in table 1. The effects of bias field on the two registration schemes was assessed by introducing a synthetic bias field into the data set. The bias fields, like the deformations, were computed by convolving a Gaussian with a random field. Each bias field was scaled and shifted so that it had a mean of 1 and variation ± 0.2 and was then applied to one of the warped images of the previous data set by pixelwise multiplication. The experiment was then repeated on the bias corrupted images. Table 2 shows the overall performance of the two algorithms in the presence of bias field. A reduction in the performance of both algorithms is apparent. The phase method, however, is significantly less affected: the increase in mean error within the brain is only 3% for the phase method, while the mean error of the intensity method increases by 14%.

5.2 Ultrasound to MR Registration

The previous examples all involve images that are related by an intensity mapping. Phase mutual information algorithms are also able to register classes of images for which this is not the case. As a practically important example of this, a set of ultrasound (US) images were simulated from deformed MR-T1 images. The T1 intensity was used to determine tissue scatter density and the intensity gradient of the T1 image was used to determine ultrasound reflectivity. These were then used to generate US-like images. Since they are generated from deformed T1 images the registration between the US images and the generating

MR images is known. An example of a simulated US image is shown in figure 4. For more details of the model of US images and registration of US images see [6].

As for the previous two experiments, a set of thirty randomly deformed ultrasound images was generated. The two algorithms were then applied with very little alteration. The phase method had to be altered slightly to ignore phase estimate too close to the edge of the US image, since the phase at these points is dominated by the image edge rather than the image contents. The only other alteration was a re-optimisation of the time step for both algorithms and an increase in the width of the regularisation Gaussian.

Figure 5 shows the distribution of registration error for the phase and intensity methods applied to the synthetic ultrasound images. In both cases, the influence of the ventricles on the registration is clearly visible. In the case of intensity MI, the improvement is limited to the area within the ventricles, while the phase MI method correctly registers some of the tissue outside the ventricles. The probable reason for this is that intensity MI is unable to deal effectively with the reflections from the ventricle walls. Since the reflections are bright they appear more likely to belong to the outside of the ventricles than the inside, which biases the intensity algorithm in favor of enlarging the ventricles. The phase algorithm is able to interpret the reflections as the superposition of an edge and a ridge and learn the correct relationship between this compound feature and the edge in the MR image.

6 Discussion

In this paper, we have described the initial development of a non-rigid registration algorithm based on local phase. To facilitate comparison with current techniques, we have analysed the extension of mutual information to use local phase. The primary motivation for this is to enable registration of certain classes of images for which intensity MI is unsuccessful. The ultrasound example (figure 5) shows that it succeeds to some extent in this.

This is not fortuitous, phase has intrinsic advantages over intensity, even for images which are well suited to registration with intensity MI. In particular,




Fig. 4. An MR-T1 image and simulated (and deformed) ultrasound image




Fig. 5. Average registration error distribution for the synthetic ultrasound images (in pixels widths) prior to registration (left), after registration with intensity MI (centre) and after registration with phase MI (right)

the low information content of its marginal distribution greatly reduces the significance of the particular choice of information measure. It is also possible to make strong assumptions about the joint distributions of phase [6]; exploiting this prior knowledge in the registration scheme is a goal of our future research.

References

1. D. Rueckert; C. Hayes; C. Studholme; P. Summers; M. Leach; D.J. Hawkes. Non-rigid registration of breast MR images using mutual information. *Proc. MICCAI*, 1998.
2. W. Crum; D.L.G. Hill; D.J. Hawkes. Information theoretic similarity measures in non-rigid registration. *Proc. IPME*, 2003.
3. J.B.A. Maintz; E.H.W. Meijering; M.A. Viergever. General multimodal elastic registration based on mutual information. *Medical Imaging*, (3338):144–154, 1998.
4. D.J. Fleet. Disparity from local weighted phase-correlation. *IEEE International Conference on SMC*, pages 48–56, October 1994.
5. M. Felsberg; G. Sommer. A new extension of linear signal processing for estimating local properties and detecting features. *Proc. DAGM*, 2000.
6. Matthew Mellor. Phase methods for non-rigid medical image registration. *D.Phil Thesis, Oxford University*, 2004.
7. Brainweb: Simulated brain database. <http://www.bic.mni.mcgill.ca/brainweb/>.