

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
30 March 2006 (30.03.2006)

PCT

(10) International Publication Number
WO 2006/032855 A2

(51) International Patent Classification:
G06T 7/00 (2006.01)

Medical Vision Laboratory, Dept of Engineering Science,
Ewert House, Ewert Place, Summertown, Oxford (GB).

(21) International Application Number:
PCT/GB2005/003592

(74) Agent: **NICHOLLS, Michael, John**; J.A. Kemp & Co.,
14 South Square, Gray's Inn, London WC1R 5JJ (GB).

(22) International Filing Date:
19 September 2005 (19.09.2005)

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0420985.4 21 September 2004 (21.09.2004) GB

(71) Applicant (*for all designated States except US*): **ISIS INNOVATION LIMITED** [GB/GB]; Ewert House, Ewert Place, Summertown, Oxford OX2 7SG (GB).


(72) Inventors; and

(75) Inventors/Applicants (*for US only*): **MELLOR, Matthew** [GB/GB]; Medical Vision Laboratory, Dept of Engineering Science, Ewert House, Ewert Place, Summertown, Oxford (GB). **BRADY, John, Michael** [GB/GB];

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,

[Continued on next page]

(54) Title: IMPROVEMENTS IN IMAGE PROCESSING



(57) Abstract: A technique for registration of images, such a multimodal images, particularly noisy images, in which the image intensities are first converted in the local energy, local phase and local orientation to provide an image descriptor based on image structure. Probability distributions and similarity calculations are carried out on the this image descriptor. The calculation of the probability distributions may include windowing with a window whose size and shape is dependent on the local energy in the image.

WO 2006/032855 A2



RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— *without international search report and to be republished upon receipt of that report*

IMPROVEMENTS IN IMAGE PROCESSING

The present invention relates to image processing and, in particular, to improvements in techniques for “registering” or aligning as accurately as possible two or more
5 images with each other.

The need to align two (or more) images arises in a variety of imaging fields. For example, in the field of medical image analysis it is often useful to align images of the same patient taken with different modalities (for example MRI, CT, PET) and/or
10 images taken at different times to monitor the development of some condition or disease. A similar requirement occurs when an image of a patient is to be compared with a standardized image-like model known as an atlas. The need for image alignment also occurs in other fields. For example, it can be useful to align aerial images of scenes taken with different sensors or frequencies bands or under different
15 imaging conditions (the sensor platform may have a different height and pose relative to the scene, and there may be cloud cover, snow cover or modifications to the scene). Image alignment may also be required in analysis of image motion: in the computation of the optic flow or the estimation of three-dimensional structure from motion. Also, in the case of multi-view image analysis such as stereo vision,
20 the simultaneously-taken images from two or more cameras need to be aligned and combined to form an estimate of the three-dimensional locations of points in the scene.

In general, image alignment techniques have proceeded by taking the two image data
25 sets, applying some form of transformation to one of them (such as a rigid transformation involving displacement and rotation, or a non-rigid transformation including, for example, warping of the structures), calculating the similarity between one of the data sets and the transformed version of the other data set, and adjusting the transformation, usually iteratively, to attempt to minimize the difference.

30
Figure 1 of the accompanying drawings illustrates the process. In step 10 the two image data sets I_1 and I_2 are acquired and in step 11 one of them is transformed using a candidate transformation. In step 12 the similarity between the transformed image and the other image is measured using a similarity measure such as cross-correlation
35 or mutual information. In practice the similarity measure is not calculated directly from the intensities, but from the individual and joint probability distributions of the intensities, and the value of the similarity measure can be calculated from these

distributions (step 12a).

For example, the mutual information of the intensities of two images, I_u & I_v , can be defined in terms of the probability distributions of the intensities of the individual pixels as follows:-

$$MI = \iint p_{uv} \ln \left(\frac{p_{uv}}{p_u p_v} \right) du dv$$

where u is an intensity value in I_u , v is an intensity value in I_v , p_u & p_v are the probabilities of those intensities occurring (i.e. the proportion of the pixels that have that intensity) and p_{uv} is the probability of two intensities occurring as a pair (i.e. at the same pixel position in both images). Other similarity measures such as normalised mutual information and correlation ratio can be similarly defined in terms of the probability distributions.

In step 13 it is checked whether the images are sufficiently aligned and, if not, the transformation is changed slightly in step 14 and steps 11 and 12 are repeated. This continues iteratively until the similarity measure has reached the required value, or, for example, has not changed significantly with the change in the transformation, at which point the process ends at step 15 with the value of the transformation which gives the best alignment having been calculated.

However, image alignment is a difficult task for several reasons. Firstly, the signal-to-noise ratio of many images is low; this is particularly true of medical images. Further, the two or more images to be aligned may have different spatial resolutions (particularly if they have been obtained by different imaging modalities), and the requirement for non-rigid transformations of one image to another make it difficult to find the correct alignment. For this reason there have been many different proposals for image alignment techniques involving different similarity criteria and different ways of coping with noise in the image.

For example, one way of allowing for noise in the image is to use a windowing technique in the calculation of the probability distribution mentioned above. Thus, whereas preparing an individual probability distribution of the intensities in an image simply involves preparing a histogram of those intensities (with each intensity value contributing only to one bin of the histogram) where a windowing technique is used, each intensity value contributes not only to its "primary" bin, but also partially to neighbouring bins. The amount of contribution to neighbouring bins depends on the

noise model chosen. For example, if a Gaussian noise model is chosen, each observed intensity value is regarded as forming the center of a Gaussian distribution, with the contribution to neighbouring histogram bins corresponding to the values of the Gaussian distribution. Different noise models can be chosen depending on the imaging modality. Thus the individual intensity values can be thought of as being “smeared-out” over several bins of the histogram. The same technique is used in the preparation of the joint probability distribution. This technique is known as Parzen windowing or kernel density estimation.

- The choice of the optimal image alignment technique, though, in any given set of circumstances, for example the choice of the correct similarity measure and the correct handling of noise, remains a difficult technical problem.

It is important to note that the term image applies not only to values of intensities calculated in two spatial dimensions, but also to image sets which are in three dimensions (in which case the pixels are usually known as voxels) and also to two or three dimensions plus time. In the present specification, though, the term “image” will be regarded as covering all such data sets, and the term “pixels” will be used to include “voxels”.

A first aspect of the present invention relates to the use of a new image descriptor for image alignment. This image descriptor consists of one or more of the local energy, local phase and local orientation in the image. The local energy, local phase and local orientation are values which can be calculated from the intensity values in an image. These are values which describe the structure in the image. One definition of energy, phase and orientation is known as the monogenic signal, as discussed, for example, in “Low- Level Image Processing with the Structure Multivector”, M. Felsberg, PhD thesis, Christian-Albrechts-Universität Kiel, 2002
<http://www.isy.liu.se/~mfe/Diss.ps.gz>.

In accordance with the present invention, rather than conducting the alignment analysis on intensity values directly, instead the analysis is conducted on an image descriptor comprising one or more of these values. Preferably the analysis is conducted at a plurality of different scales. Each of local energy, local phase and local orientation presuppose analysis of the image at a certain scale. Images may be aligned more quickly by starting the analysis at a relatively large scale (this may correspond, for example, to aligning the overall outline of two image structures),

followed by analysis at successively smaller scales. It should be noted, as mentioned above, that the scales may be purely spatial or temporal.

A second aspect of the invention relates to an adaptation of the way in which the probability distributions of the image descriptor are calculated in order better to
5 account for the noise in the image. In particular, a windowing function is used when calculating the probability distributions and the size and shape of the window are varied in accordance with the local structure in the image. In particular, the size and shape of the window can be varied in accordance with the local energy as calculated
10 at each position in the image. Preferably the greater the local energy the smaller the window. This reduces the susceptibility of the joint probability distribution estimate to the effects of noise. Furthermore, because local energy is concentrated around feature points in the image, the method becomes feature driven (ie. features in the image are used preferentially in the alignment process as the contribution to the
15 histograms in these areas is less "smeared out").

Preferably the aspects are combined together so that the image descriptor consisting of one or more of the local energy, local phase and local orientation is used as the image descriptor in the calculation of the probability distributions using an energy-
20 dependent window as mentioned above .

The invention will be further described by way of example with reference to the accompanying drawings in which:-

25 Figure 1 is a flow diagram illustrating the general technique of image alignment used in the prior art;

Figure 2 is a flow diagram illustrating the general image alignment method in accordance with an embodiment of the present invention;

30

Figure 3 illustrates schematically the calculation of the image descriptor in accordance with one embodiment of the present invention;

Figure 4 illustrates schematically the calculation of the probability
35 distribution of image descriptors in accordance with an embodiment of the present invention;

Figures 5A to I are flow diagrams and pseudo-code schematically illustrating

an example of the present invention; and

Figures 6 (a) to (d) illustrate an example of image alignment using an embodiment of the present invention.

5

Referring first to Figure 2 the general processing flow of image alignment using an embodiment of the present invention will be described. Firstly, in step 20, the two sets of image data to be aligned or "registered" are received. As discussed above, these may be two-dimensional images or three-dimensional images, and may be
10 supplemented by time. In step 21 one or more scales are chosen for the image analysis. Preferably a multi-scale approach is used, with the analysis first being conducted at a relatively large scale, and then successively smaller scales (corresponding to finer resolutions). In the example which is illustrated in Figure 6, four different scales were chosen. The scales are chosen based on the size of the
15 structures in the image and the accuracy of the registration required.

In step 22 the monogenic signal is calculated for each image point at the chosen scale. The monogenic signal consists of the local energy, the local phase and the local orientation. Methods of calculating the monogenic signal are known, and any
20 of the available techniques may be used with the present invention, for example steerable quadrature filters such as described in "The Design and Use of Steerable Filters" by Freeman and Adelson (PAMI 13(9), Sept 1991, 891-906) or by the techniques described in "Low- Level Image Processing with the Structure Multivector", M. Felsberg, PhD thesis, Christian-Albrechts-Universität Kiel, 2002
25 <http://www.isy.liu.se/~mfe/Diss.ps.gz>. This is discussed in more detail below.

In step 23 the candidate transformation is applied to one of the data sets and in step 24 the probability distributions needed for calculation of the similarity measure in step 25 are calculated. As will be explained below, in accordance with an aspect of
30 the present invention, the probability distributions are calculated taking into account the noise in the image in an adaptive way in accordance with the local structure in the image. Further, they are calculated on the components of the monogenic signal - the local energy, local phase and local orientation, not on the intensities as in the prior art. Step 26 and 27 correspond to the normal iterative processing to optimize
35 the transformation and in step 28 it is checked whether all scales have been completed. If not then steps 22 to 26 are repeated for each of the chosen scale until they are all completed, at which point the process ends at step 30 with best estimate

for the transformation having been calculated.

The local energy, local phase and local orientation, are indicative of the local properties of a region within an image. For example, they describe such features as step changes, ridges, sudden spikes and so on. Image features can be classified according to their local symmetry. Just as any real signal can be locally decomposed into a symmetric part and anti-symmetric part (corresponding, for example, to the sine and cosine Fourier components), the same is true of image features. For example, an edge between a dark area and a light area is locally anti-symmetric while a peak in intensity is locally symmetric. The ratio of the amplitude of the symmetric and anti-symmetric components in any given feature amounts to a simple summary of the local shape of the intensity profile of the feature. The combined amplitude gives information about the amount of local signal activity, and reaches a local maximum near the center of the feature. These two quantities are what is measured by the phase and energy in the monogenic signal respectively. If considering, therefore, the variation of intensity across an image feature such as an edge from white to black, this feature is locally anti-symmetric and so could be described as having mainly a sine component and thus a phase close to say, zero. On the other hand, a feature which is symmetric (and thus over the same distance goes from white to black and white again) would have predominantly a cosine component and thus a phase of nearer ninety degrees. The phase and energy can be obtained from an image data set by convolving it with a pair of quadrature filters, but with a two (or more) dimensional image there is the problem that the image features are different in different directions across the image. Thus the monogenic signal also includes the feature orientation which can intuitively be understood as the direction associated with a local linear structure (such as the direction of an edge). The three quantities of local phase, local energy and local orientation can be obtained by using three filters which are convolved with the image, using the known techniques described in the references above.

30

Figure 3 illustrates this schematically in which two odd filters, one for the vertical direction in the image and one for the horizontal direction, and one isotropic even filter are used and their outputs are combined to calculate the three components of the monogenic signal: the local energy A_s , the local phase ϕ and the local orientation θ .

35

Thus as a result of convolving each of the image data sets to be aligned with the three

filters, the intensity values are converted into values at each image point (pixel) for the components of the monogenic signal - the local energy, the local phase and the local orientation.

- 5 The next step involves calculating the individual probability distributions in the two data sets, and also the joint probability distribution corresponding to steps 23 and 24 of Figure 2.

- As explained above a conventional way of estimating a probability distribution
- 10 taking account of noise is to use a windowing function, such as a Parzen window in the construction of the histogram or joint histogram. In this approach each value being added to the histogram contributes not to a single bin, but contributes over several bins in a distribution which is based on the noise model for the data set. Thus considering a signal I corrupted by additive white Gaussian noise, the distribution
 - 15 would be:-

$$P(I|I_k) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(I-I_k)^2}{2\sigma^2}} \quad \text{Equation 1}$$

where σ is the standard deviation of the noise.

- 20 Thus this equation tells you how much of the vote of observed signal I_k goes into each bin I of the histogram. Conventionally this equation is used to calculate the contribution at each image point in the same way. However, with the present invention, this windowing function is adjusted at each point in accordance with the local image structure. In this embodiment it is adjusted by the local energy A_k at that
- 25 point so that, for example, taking the phase component of the monogenic signal, the Parzen window is defined as:-

$$P(\phi|\phi_k, \sigma, A_k) = \frac{A_k}{\sigma\sqrt{2\pi}} e^{-\frac{A_k^2(\phi-\phi_k)^2}{2\sigma^2}} \quad \text{Equation 2}$$

30

Thus if the observed phase is ϕ_k , the histogram bin for phase value ϕ receives the value calculated by equation 2 where A_s is the local energy (also calculated from the monogenic signal), and σ is the standard deviation of the noise model, which is estimated for each image data set, for example by modelling the imaging process or

by calculating the variance of the intensity in the image. So the amount of
 “smearing” of the observed values represented by the windowing function varies
 from point to point in the image with the local energy: In essence, σ , which governs
 the size and shape of the window, is modified by the local energy, which is an aspect
 5 of image structure.

The equivalent window expressions for the local orientation and local energy
 respectively in each of the two data sets are:-

$$P(\theta | \theta_k, \sigma, A_o) = \frac{A_o}{\sigma\sqrt{2\pi}} e^{\frac{-A_o^2(\theta - \theta_k)^2}{2\sigma^2}}$$

$$P(A | A_k, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(A - A_k)^2}{2\sigma^2}}$$

where A_o is the odd component of the energy, calculated as the square root of the
 sum of the two odd parts of the monogenic signal at each scale.

Figure 4 illustrates schematically how the probability distributions for the
 15 components of the monogenic signal are calculated. A linear noise model 40 (for
 example the Gaussian noise model discussed above or any other noise model suitable
 for the image data set) is convolved with the same filters 41, 42 and 43 to generate
 the monogenic signal and is then utilized in the estimation at 44, 45 and 46 of the
 probability distribution functions 44, 45 and 46 for the local energy, local orientation
 20 and local phase.

These three PDFs together form the PDF of the new variable which will be used in
 place of intensity in the calculation of the similarity between the image data sets.

25 These probability distributions can then be used to calculate the similarity between
 the two images, using any desired similarity measure as illustrated in step 25 of
 Figure 2.

Figure 6 illustrates an example of two test images to which the present invention was
 30 applied. Figure 6a and Figure 6b illustrate two significantly different images of the
 same scene. Figure 6c illustrates the most likely horizontal alignment of the two
 images as calculated by the use of traditional mutual information of the intensities in
 the two image data sets. It can be seen that there is a peak at zero, implying that the

images should be aligned at this point. Figure 6d illustrates the correct alignment point as calculated using an embodiment of the present invention. In particular the method of Figure 2 was run at four different scales and assuming a simple white, additive Gaussian noise model. It can be seen that again there is a peak at zero, illustrating that the correct alignment was calculated, but that this peak is much clearer than the peak in Figure 4c.

A simple example of a complete PDF estimation and registration procedure in accordance with an embodiment of the invention will now be described with reference to Figures 5A to 5I. Figure 5A illustrates the overall process. First, the images are received in step 50 and approximately aligned (e.g. by rigid registration algorithm as known in the art (not illustrated in Figure 5A)). One image is designated the source image and the other the target image. The source image will be deformed until it is aligned with the target image. Local energy, local phase and local orientation are then calculated at step 51 (in process A which will be described in more detail below) for each image, at a particular scale, which may be chosen by the designer to suit the problem at hand.

Next, in step 52 the initial joint Probability Distribution (PDF) of the descriptors must be estimated, for example using process B as shown by the pseudocode in Figure 5F. Then in step 53 the incremental improvement to the transformation is calculated, for example by process C shown in Figures 5G to 5I and the new transformation is applied to the source image. This process is continued iteratively until the improvement DF in alignment caused by the change to the transformation is less than a threshold lim as tested at step 55, finishing at step 56 with the best alignment achieved.

In process A local energy, local phase and local orientation are calculated from the response of a set of three filters at each scale, as shown in figures 5C, 5D and 5E. The choice of the filter family is application specific; in this example a family of non-blurring filters are used. The overall flow is shown in Figure 5B. For each of N scales, where N is five in this example, a set of three filters is constructed in step A1. First the rotationally symmetric filter is defined. As shown in Figure 5C, in this example, the rotationally symmetric filter is a weighted sum of two other rotationally symmetric filters, A and B , defined by $A=r^{2.5+0.5*scale}$, $B=r^{2+0.5*scale}$, where r is the distance from the filter centre, calculated in step $r=sqrt(x^2+y^2)$, wherein the function $sqrt()$ performs a square root operation. The weight applied to each filters is defined

as one over the summation over all x and y of the filter value, as shown in the steps $A=A/\text{sum}(A)$ and $B=B/\text{sum}(B)$, where the function $\text{sum}()$ performs the summation of x and y . Next, each image is convolved with the rotationally symmetric filter formed by subtracting B from A , as in the steps $\text{Bandpass_1}_{\text{scale}}=\text{convolve}(\text{Image_1}, A-B)$ and
 5 $\text{Bandpass_2}_{\text{scale}}=\text{convolve}(\text{Image_2}, A-B)$. This is repeated for each of the $N=5$ scales, which results in a set of band-pass filtered images.

Next, as shown in figure 5D, the odd filter responses are calculated from the band-pass filtered images. This is achieved by applying the anti-symmetric monogenic
 10 filters $H1$ and $H2$ described in "Low- Level Image Processing with the Structure Multivector", M. Felsberg, PhD thesis, Christian-Albrechts-Universität Kiel, 2002 <http://www.isy.liu.se/~mfe/Diss.ps.gz>. These are constructed in the steps $H1=x/(2\pi(x^2+y^2)^{3/2})$ and $H2=y/(2\pi(x^2+y^2)^{3/2})$.

15 Next for each of the $N=5$ scales, each of the bandpass filters is convolved with each of $H1$ and $H2$ to form the odd filtered images, as in step
 $\text{Odd_1}_{\text{scale}}=\text{convolve}(\text{Bandpass_1}_{\text{scale}}, H1)$ and the three subsequent steps. Finally, the local energy and local phase may be calculated, as in figure 5E.

20 If desired, local orientation could also be calculated at this point, according to the method described in "Low- Level Image Processing with the Structure Multivector", M. Felsberg, PhD thesis, Christian-Albrechts-Universität Kiel, 2002. For each of the $N=5$ scales, the odd and even components are calculated. In the present example, the even component, at each scale and for each image, is simply the corresponding band-pass filtered image, as in step $\text{Even_1} = \text{Bandpass_1}_{\text{scale}}$ and the subsequent step. The
 25 odd component is calculated, at each scale and for each image, by taking the square root of the sum of the two corresponding odd filtered images squared, as in step
 $\text{Odd_1} = (\text{Odd_1_x}_{\text{scale}}^2 + \text{Odd_1_y}_{\text{scale}}^2)^{1/2}$ and the subsequent step. The local phase at each scale and for each image is then calculated according to the step
 30 $\phi_{\text{1scale}}=\arctan(\text{Even_1}/\text{Odd_1})$ and the subsequent step, where the $\arctan()$ function calculates the arctangent. The local energy is calculated by taking the square root of the sum of the odd and even components squared, as in the step
 $A_{\text{1scale}}=(\text{Even_1}^2 + \text{Odd_1}^2)^{1/2}$ and the subsequent step. This completes the estimation of local phase and local energy in every point in each image, at each scale.

35

The next step in the registration procedure is to estimate the joint PDF, as shown in figure 5F. In this simple example, the descriptor will consist of only the local phase.

The first step in estimating the joint PDF is to define an image noise model; typically this might be Additive White Gaussian distributed Noise (AWGN). This only needs to be defined once, at the beginning of the process, and may be obtained, for example, from a prior knowledge of the noise characteristics of the imaging apparatus. Once the noise model is defined, the PDF is initialised. In the present example, the PDF is represented by a two dimensional histogram of ten thousand bins (one hundred bins square), initialised to zero. PDF initialisation is performed in the steps $Bin_no = 100$ and $Joint_hist_{scale} = \text{zeros}([Bin_no, Bin_no])$, where the function $\text{zeros}(u, v)$ creates a matrix of size u by v , populated with zeros.

10

The next step is construct the PDF for the observed phase values, defined by Equation 2, using the observed value of the local energy as A_k and the standard deviation of the AWGN model as σ . This is done once for each of the five scales, within the loop $For\ scale = 1:5$. To construct the PDF at the present scale, each pixel is visited in turn, within the loop $For\ x=1:x_size, y=1:y_size$, where x_size and y_size are the sizes of the images in the x and y directions respectively. For each pixel, the PDF indices P1 and P2, corresponding to the observed phases at the present pixel in each image, and their associated local energies A1 and A2 are retrieved, as in step $P1 = \phi_{I_{scale}}(x, y)$ and the three subsequent steps. These values are then used to increment each bin of the PDF in turn, in the loop $For\ I=1:Bin_no, j=1:Bin_no$. Within this loop, each bin of the PDF is incremented with a number proportional to the joint probability of the true phase lying within that bin, according to the noise model in equation 2. This achieved by the line $Joint_hist_{scale}(i, j) += \exp(-(A1^2(P1-i)^2 + A2^2(P2-j)^2)/(2*\sigma^2))$, where i and j are the PDF indices and σ is the standard deviation of the image noise. Once each bin of the PDF has been incremented once for each pixel pair, the PDF is normalised so that it sums to one, as in the step $Joint_hist_{scale} = Joint_hist_{scale}/\text{sum}(Joint_hist_{scale})$. The final step is to calculate the first and second marginal distributions by summing the PDF along for all values of first and second phase respectively, which is accomplished by the final two lines in the outer loop of figure 5F.

From the PDF it is possible to refine the registration using a similarity measure, such as mutual information, and one of the many prior art methods for minimising it. This example employs a simplified variant of the method described in Information Theoretic Similarity Measures in Non-rigid Registration, William R. Crum, Derek L. G. Hill, David J. Hawkes, in Information Processing and Medical Imaging, 2003 pp 378-387. This algorithm is illustrated in figure 5g. Each scale is treated separately.

35

The first step for each scale is to look up the joint and marginal probabilities of the observed phases at each pixel to create three probability maps. This is performed by the code in figure 5h. For each pixel, the two observed phases are translated into PDF indices P1 and P2 by the *step* $P1 = 1 + \text{round}(198 * \phi_{scale}(x,y) / \pi)$ and the subsequent
5 step. In these two steps, the function *round()* rounds its argument to the nearest integer value and the factors 198 and pi are chosen to guarantee that the resulting index value lies somewhere in the interval 1 to *Bin_no*. These indices are then used to look up the joint and marginal probabilities at the present point, as in the last three steps in the loop of figure 5H.

10

When this procedure has been done for every pixel, at the present scale, increments to the x and y components of the registration *Flow_x_Increment* and *Flow_y_Increment* are calculated according to the method in Information Theoretic Similarity Measures in Non-rigid Registration, William R. Crum, Derek L. G. Hill, David J. Hawkes, in Information Processing and Medical Imaging, 2003 pp 378-387.
15 This is performed by the code in figure 5i. Once this whole procedure has been completed for every pixel at every scale, the resulting registration update must be regularised. This could be achieved by any of the extensive set of methods known in the literature. The optimal choice will depend on the application; convolution with a
20 Gaussian window function is an example of a simple and effective regularisation method, which works well in the present implementation. The incremental improvement is then added to the previous estimate of the registration, and the source image is deformed accordingly, using one of the prior art interpolation schemes. The widely used bilinear interpolation scheme has been found to work well
25 in the present implementation.

One iteration of registration algorithm, as shown in Figure 5A, is now complete. The mean distance moved by the incremental improvement is measured and this it is above a certain (user defined) limit, the whole process is repeated, starting with the
30 re-computation of local energies and phases. When an iteration occurs in which the incremental change is smaller than the user defined limit (step 55), the procedure terminates and alignment is achieved (step 56).

CLAIMS

1. A method of measuring the similarity between two medical image data sets, said data sets
5 comprising intensity values, comprising calculating from the intensity values the values for each image data set of an image descriptor comprising at least one of local energy, local phase and local orientation, and using the calculated an image descriptor in the calculation of a similarity function to give a measurement of similarity between the two image data sets.
- 10 2. A method according to claim 1 wherein the image descriptor comprises at least two of local energy, local phase and local orientation.
3. A method according to claim 1 wherein the image descriptor comprises at least three of local energy, local phase and local orientation.
- 15 4. A method according to claim 1, 2 or 3 wherein the values of the image descriptor are calculated at a plurality of different scales.
5. A method according to claim 4 wherein said different scales are different
20 spatial scales.
6. A method according to claim 4 or 5 wherein said different scales are different temporal scales.
- 25 7. A method according to any one of claims 1 to 6 wherein similarity measure is selected from Mutual Information, Joint Entropy, Correlation Ratio, Normalised Mutual Information, Woods Ratio, Entropy Correlation Coefficient.
8. A method of analysing an image, the method comprising calculating the
30 probability distribution of the values of data set comprising values of an image descriptor, wherein a windowing function is applied to said data set, said windowing function having a window, wherein at least one of the size and shape of the window is varied through the data set in dependence upon the local structure of the image.
- 35 9. A method determining the similarity between two image data sets by calculating the value of a similarity function from the probability distributions of the value of an image descriptor in each of the two image data sets, wherein calculation of the probability distributions

includes the step of applying a windowing function to said data set, said windowing function having a window, wherein at least one of the size and shape of the window is varied through the data set in dependence upon the local structure of the image.

5 10. A method according to claim 8 or 9 wherein the windowing function is based on a structure dependent model of the noise in the local descriptors .

11. A method according to claim 8, 9 or 10 wherein at least one of the size and shape of the window is dependent on the local energy in the image.

10

12. A method according to claim 8, 9 or 10 wherein at least one of the size and shape of the window is dependent on the local phase in the image.

15

13. A method according to claim 8, 9 or 10 wherein at least one of the size and shape of the window is dependent on the local orientation in the image.

14. A method according to claim 8, 9 or 10 wherein at least one of the size and shape of the window is dependent on at least two of the local energy, the local phase and the local orientation in the image.

20

15. A method according to claim 8, 9 or 10 wherein at least one of the size and shape of the window is dependent on at least three of the local energy, the local phase and the local orientation in the image.

25

16. A method according to any one of claims 8 to 15 wherein the image descriptor comprises at least one of local energy, local phase and local orientation.

17. A method according to any one of claims 8 to 15 wherein the image descriptor comprises at least two of local energy, local phase and local orientation.

30

18. A method according to any one of claims 8 to 15 wherein the image descriptor comprises all three of local energy, local phase and local orientation.

19. A method according to any one of claims 8 to 18 wherein the image descriptor is
35 calculated at a plurality of different scales.

20. A method according to claim 19 wherein said different scales are different spatial scales.


21. A method according to claim 19 or 20 wherein said different scales are different temporal scales.

5 22. A computer program comprising program code means for executing on a programmed computer the method of any one of the preceding claims.

23. An image processing system for processing images in accordance with the method of any one of the preceding claims.

1/13

Fig.1.



2/13

Fig.2.






Fig.3.




4/13

Fig.4.




5/13

Fig.5A.



6/13

Fig.5B.



7/13

Fig.5C.

A1

$$r = \sqrt{x^2 + y^2}$$

For Scale = 1:N{

$$A = r^{2.5 + 0.5 * \text{Scale}}$$

$$B = r^{2 + 0.5 * \text{Scale}}$$

$$A = A / \text{sum}(A)$$

$$B = B / \text{sum}(B)$$

$$\text{Bandpass_1_scale} = \text{convolve}(\text{Image_1}, A - B)$$

$$\text{Bandpass_2_scale} = \text{convolve}(\text{Image_2}, A - B)$$

}end

Fig.5D.

A2

$$H1 = x / (2\pi(x^2 + y^2)^{3/2})$$

$$H2 = y / (2\pi(x^2 + y^2)^{3/2})$$

For Scale = 1:N{

$$\text{Odd_1_x_scale} = \text{convolve}(\text{Bandpass_1_scale}, H1)$$

$$\text{Odd_1_y_scale} = \text{convolve}(\text{Bandpass_2_scale}, H2)$$

$$\text{Odd_2_x_scale} = \text{convolve}(\text{Bandpass_2_scale}, H1)$$

$$\text{Odd_2_y_scale} = \text{convolve}(\text{Bandpass_2_scale}, H2)$$

}end

8/13

Fig.5E.

A3

For Scale = 1:N{

Even_1 = Bandpass_1_{scale}

Even_2 = Bandpass_2_{scale}

Odd_1 = (Odd_1_x²_{scale} + Odd_1_y²_{scale})^{1/2}

Odd_2 = (Odd_2_x²_{scale} + Odd_2_y²_{scale})^{1/2}

$\phi_{1scale} = \arctan(\text{Even}_1 / \text{Odd}_1)$

$\phi_{2scale} = \arctan(\text{Even}_2 / \text{Odd}_2)$

$A_{1scale} = (\text{Even}_1^2 + \text{Odd}_1^2)^{1/2}$

$A_{2scale} = (\text{Even}_2^2 + \text{Odd}_2^2)^{1/2}$

}end

9/13

Fig.5F.

B

Bin_no = 100

For scale = 1:N{

Joint_hist_{scale} = zeros([Bin_no,Bin_no])Marginal_1_{scale} = zeros([Bin_no,1])Marginal_2_{scale} = zeros([Bin_no,1])

For x = 1:x_size, y=1:y_size{

P1 = ϕ_1 _{scale}(x,y)P2 = ϕ_2 _{scale}(x,y)A1 = A_1_{scale}(x,y)A2 = A_2_{scale}(x,y)

For i=1:Bin_no, j=1:Bin_no,

Joint_hist_{scale}(i,j) +=exp(-(A1²(P1-i)²+A2²(P2-j)²)/(2* σ^2))

}end


}end

Joint_hist_{scale}=Joint_hist_{scale}/sum(Joint_hist_{scale})Marginal_1_{scale}(P1) =sum(Joint_hist_{scale},1)Marginal_2_{scale}(P2) =sum(Joint_hist_{scale},2)

}end

10/13

Fig.5G.



11/13

Fig.5H.

C1

```

For x = 1:x_size, y=1:y_size{

    P1 = 1+round(198* $\phi_1$ _scale(x,y)/ $\pi$ )
    P2 = 1+round(198* $\phi_2$ _scale(x,y)/ $\pi$ )

    Joint_probability_map(x,y) = Joint_hist_scale(P1,P2)
    Marginal_map_1(x,y) = Marginal_1_scale(P1)
    Marginal_map_2(x,y) = Marginal_2_scale(P2)

}end

```

Fig.5I.

C2

```

For x = 1:x_size, y=1:y_size{

    Pjr = Joint_probability_map(x+1,y)
    Pjl = Joint_probability_map(x-1,y)
    Pju = Joint_probability_map(x,y+1)
    Pjd = Joint_probability_map(x,y-1)

    Pr = Marginal_map_1(x+1,y)
    Pl = Marginal_map_1(x-1,y)
    Pu = Marginal_map_1(x,y+1)
    Pd = Marginal_map_1(x,y-1)

    Flow_y_Increment(x,y) = [ log(Pjd./Pd)+...
                               Log(Pju./Pu) ]

    Flow_x_Increment(x,y) = [ log(Pjl./Pl)+...
                               Log(Pjr./Pr) ]

}end

```

12/13

Fig.6A.





Fig.6B.



13/13

Fig.6C.




Fig.6D.

