Incorporating Connected Region Labelling into Automated Image Registration Using Mutual Information

C.Studholme, D.L.G.Hill, D.J.Hawkes
Radiological Sciences
UMDS, Guys Hospital
London SE1 9RT
{cs,dlgh,djh}@umds.ac.uk

Abstract

The information theoretic measure of mutual information has been successfully applied to multi-modality medical image registration for several applications. There remain however, modality combinations for which mutual information derived from the occurrence of image intensities alone does not provide a distinct optimum at true registration. We propose an extension of the technique through the use of an additional information channel supplying region labelling information. These labels, which can specify simple regional connectivity or express higher level anatomical knowledge, can be derived from the images being registered. We show how the mutual information measure can be extended to include an additional channel of region labelling, and demonstrate the effectiveness of this technique for the registration of MR and PET images of the pelvis.

Keywords: Automated Image Registration, Voxel Similarity Measures, Region Labelling, Mutual Information.

1 Introduction

It has been shown that optimisation of voxel similarity based registration measures [5], in particular that of mutual information [8, 1, 4], provides a robust approach to the registration of brain images acquired with MR and PET, or MR and CT.

Experimentally we have found problems can occur when using mutual information to align other modality combinations or images acquired in other regions of the body. In some cases this may be due to a of a lack in differentiation between spatially unconnected regions in one modality which are connected in the other modality. This is a particular problem when the two images are truncated or when the corresponding structures in the images essential in defining the registration are significantly perturbed by noise. Optimisation of mutual information, which can be derived from the joint and separate probability distributions of image intensities, contains no information on either the spatial distribution of intensities or the connectivity of uniform regions.

In principle a complete segmentation of the MR image, in which each distinguishable anatomical structure has a unique label, could be registered by maximising mutual information between the labelled image and the lower contrast, lower resolution image. In practice fully automated, complete and accurate regional labelling is impossible in anatomically complex and noisy images. It is difficult to ensure connectivity of regions of the same anatomical structure or tissue type while maintaining separation of regions corresponding to different structures or tissue types. In particular the identification of truly corresponding anatomical boundaries in different modalities is non-trivial.

In previous work mutual information between two channels, derived from the intensities in a each modality, has been maximised to achieve registration. We propose the introduction of further channels containing information in the form of labelling of distinct regions within the images. An additional channel of region labels allows differentiation of unconnected regions of similar intensity within an image.

In this paper we introduce a formalism for a similarity measure enabling the inclusion of coarse labelling of regions within the high resolution, high contrast image. An information theoretic approach enables the original image and region labelling to be combined in the registration process making use of information provided by both. We discuss how a simple region labelling process introduces additional information into the registration process. The behaviour of the approach is demonstrated in the registration of MR and PET images of the pelvis, conventionally difficult to register, by the addition of a coarsely labelled anatomical map derived from the MR image.

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31,
IEEE Computer Society Press, ISBN 0-8186-7367-2

2 Mutual Information

2.1 Two Images

Given a pair of images to register and a transformation mapping one set of voxels onto the other, we can find for corresponding voxels, the intensities $m \in M$ in image $\mathbf{m}$, and intensities $n \in N$ in a lower contrast, lower resolution image $\mathbf{n}$. In the example considered in this paper $\mathbf{m}$ is an MR image of the pelvis and $\mathbf{n}$ is a nuclear medicine PET image. The sets M and N therefore depend on the rigid body transformation between the images. We can calculate the probability of occurrence of individual intensities $p\{m\}$ from the image $\mathbf{m}$, and intensities $p\{n\}$ from the image $\mathbf{n}$, and also the probabilities of corresponding intensity pairs $p\{m, n\}$.

The mutual information measure is derived from a statistical analysis of a communication channel, and is a measure of corresponding or mutual information between the transmitted and received signals [3]. This measure is alternatively referred to as relative entropy or transinformation and makes no assumption about the functional form of the channel. Mutual information has been proposed independently for various medical image registration applications by Collignon et al [1] and Viola and Wells [7] [8]. The mutual information between an image $\mathbf{m}$ with intensity m and an image $\mathbf{n}$ with intensity n defined from the 2D probability distribution is:

$$I(M; N) = \sum_{m \in M} \sum_{n \in N} p\{m, n\} \log \frac{p\{m, n\}}{p\{m\}p\{n\}} \quad (1)$$

This can be expressed in terms of the information present in image $\mathbf{m}$ $H(M)$, the image $\mathbf{n}$ $H(N)$, and the combined image $H(M, N)$:

$$I(M; N) = H(M) + H(N) - H(M, N) \quad (2)$$

In other words by maximising mutual information we minimise the information in the combined image with respect to that present in the two component images. If there are any shared features in the two images, at misregistration they are duplicated in the combined image. This duplication of features results in an excess of information in the combined image with respect to that provided by the two images separately. The relationship between the information provided by the images can be illustrated in a set theory form as shown in figure 1. The sets on the left illustrate the entropies of the two images separately. The central diagram represents the joint entropy of the combined image when a transformation maps the intensities of one onto the intensities of the other. The right hand figure shows that the mutual information is the replicated information from the two images.

2.2 Three Images

In our application we introduce further information in the form of a labelling of image $\mathbf{m}$ prior to registration, by creating a third image $\mathbf{l}$ with labels $l \in L$. This is shown graphically in figure 2. The measure of mutual information can be directly extended for this case as the difference between the sum of the information in the three images separately, and the information in the combined image:

$$I(M; L; N) = H(M) + H(L) + H(N) - H(M, L, N) \quad (3)$$

Images $\mathbf{m}$ and $\mathbf{l}$ are inherently registered and we are interested in the information shared between the PET image and these images which is illustrated by the set representation in figure 3.

This can be written algebraically as:

$$I(M, L; N) = H(M, L) + H(N) - H(M, L, N) \quad (4)$$

The entropies $H(M, L)$ $H(N)$ $H(M, L, N)$ can be derived from the probability of occurrence of intensities M and N, and labels L:

$$H(M, L) = \sum_{m \in M} \sum_{l \in L} p\{m, l\} \log \frac{1}{p\{m, l\}} \quad (5)$$

$$H(N) = \sum_{n \in N} p\{n\} \log \frac{1}{p\{n\}} \quad (6)$$

$$H(M, L, N) = \sum_{m \in M} \sum_{l \in L} \sum_{n \in N} p\{m, l, n\} \log \frac{1}{p\{m, l, n\}} \quad (7)$$

and substituted into equation 4 giving:

$$I(M, L; N) = \sum_{m \in M} \sum_{l \in L} p\{m, l\} \log \frac{1}{p\{m, l\}} + \sum_{n \in N} p\{n\} \log \frac{1}{p\{n\}} - \sum_{m \in M} \sum_{l \in L} \sum_{n \in N} p\{m, l, n\} \log \frac{1}{p\{m, l, n\}}$$

Given that $p\{m, l\} = \sum_{n \in N} p\{m, l, n\}$ and $p\{n\} = \sum_{m \in M} \sum_{l \in L} p\{m, l, n\}$,

$$I(M, L; N) = \sum_{m \in M} \sum_{l \in L} \sum_{n \in N} p\{m, l, n\} \log \frac{1}{p\{m, l\}} + \sum_{n \in N} \sum_{m \in M} \sum_{l \in L} p\{m, l, n\} \log \frac{1}{p\{n\}} - \sum_{m \in M} \sum_{l \in L} \sum_{n \in N} p\{m, l, n\} \log \frac{1}{p\{m, l, n\}}$$

this can be rearranged to give a simpler form:

$$I(M, L; N) = \sum_{m \in M} \sum_{l \in L} \sum_{n \in N} p\{m, l, n\} \log \frac{p\{m, l, n\}}{p\{m, l\}p\{n\}} \quad (8)$$

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31, IEEE Computer Society Press, ISBN 0-8186-7367-2

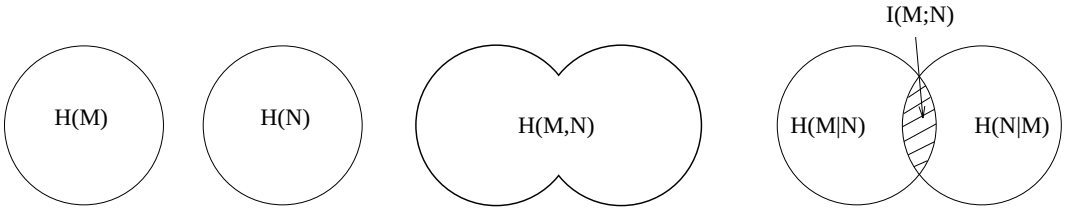

Figure 1. A set theory representation of the entropies involved when combining two images

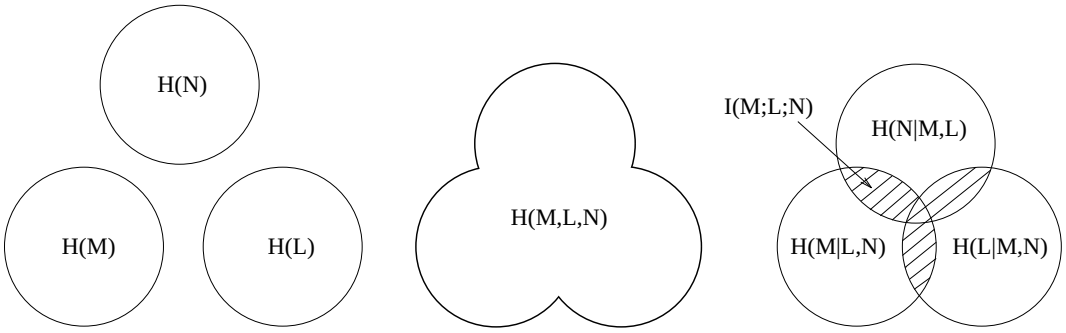

Figure 2. A set theory representation of the entropies involved when combining three images

2.3 Information and Region Labelling

In order to illustrate how extra regional information is being introduced we consider a simple region extraction process. This may consist of a thresholding to distinguish boundaries of interest, followed by a connected component labelling of the identified regions as illustrated in figure 4. Thresholding is effectively combining intensities (symbols) in the image, producing only a subset of the information in the original image. If a binary threshold is used then the image is being reduced to two symbols $T = \{t_0, t_1\}$ where, if there are i intensities in the original image and m_t is the threshold intensity so,

$$t_0 = \{m_0, \ldots m_t\} \quad (9)$$

$$t_1 = \{m_{t+1}, \ldots m_i\}. \quad (10)$$

The information contained in this thresholded image is then given by,

$$H(T) = -\sum_{t \in \{t_0, t_1\}} p\{t\} \log p\{t\}. \quad (11)$$

Using the additivity property of information measures (see Reza[3], page 84) it can be shown that the information content of a set of symbols cannot be decreased when symbols are partitioned so,

$$H(m_0, \ldots m_t, m_{t+1} \ldots m_i) \geq H(t_0, m_{t+1}, \ldots m_i) \quad (12)$$

From this we can say that the information content of the thresholded image is always less than or equal to that in the original image,

$$H(T) \leq H(M). \quad (13)$$

In addition the joint entropy between the image and its thresholded version is given by,

$$H(M,T) = -\sum_{t \in T} \sum_{m \in M} p\{m,t\} \log p\{m,t\} \quad (14)$$

which can be re-written in terms of the occurrence of the original intensities only,

$$H(M,T) = -\sum_{m \in t_0} p\{m\} \log p\{m\} - \sum_{m \in t_1} p\{m\} \log p\{m\}, \quad (15)$$

giving $H(M,T) = H(M)$. Using this, the mutual information between the image and its threshold,

$$I(M;T) = H(T) + H(M) - H(M,T). \quad (16)$$

becomes simply $I(M;T) = H(T)$, i.e. the information in the thresholded image $H(T)$ is simply a subset of that provided by the original image $H(M)$.

Labelling unconnected regions differently, effectively repartitions the threshold labels and conversely can only maintain or increase the number of symbols (intensities or labels) present in the image. Applying the additivity property again, this must always maintain or increase the amount of information so that, if the threshold label t_0 is being divided up into unconnected regions $\{l_0 \ldots l_k\}$,

$$H(l_0 \ldots l_k, t_1) \geq H(t_0, t_1). \quad (17)$$

If $\{l_0, \ldots l_k\}$ is simply a repartitioning of $\{m_0, \ldots m_t\}$ (i.e. no regions with the same intensity are spatially unconnected) then $H(M,L) = H(M)$ and so $I(M,L;N) =$

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31, IEEE Computer Society Press, ISBN 0-8186-7367-2

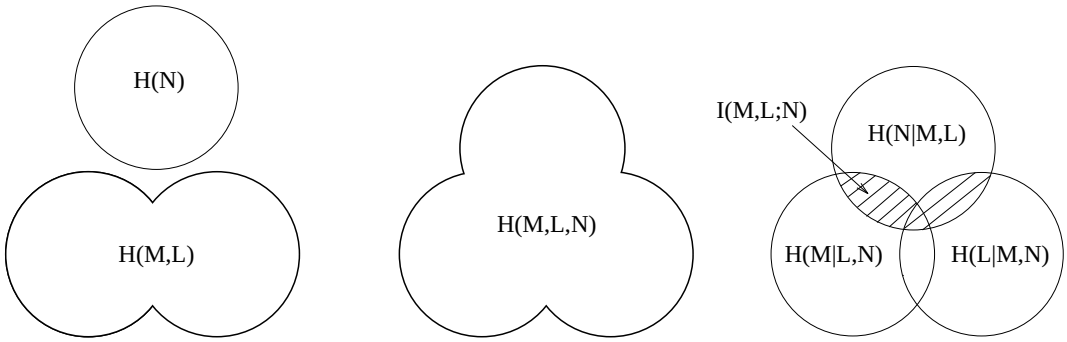

Figure 3. A set theory representation of the entropies involved when combining two images with a third

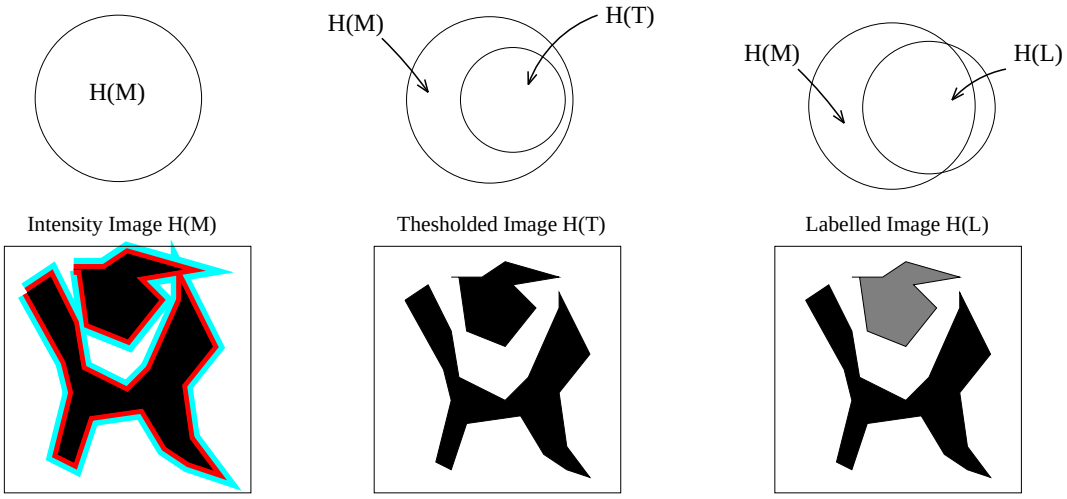

Figure 4. Simple image region labelling consisting of thresholding and connected component labelling.

$I(M;N)$. If though, the image contains cases where an intensity $m \in t_0$ is partitioned into more than one label l then,

$$H(M,L) \geq H(M) \tag{18}$$

and similarly $H(M,L,N) \geq H(M,N)$, so from figure 3,

$$I(M,L;N) \geq I(M;N). \tag{19}$$

By introducing connected region labelling information we can only maintain or increase the mutual information in the registration process. What we would like is a labelling for which $I(M;L)$ is minimised (there are no duplicated features in the labelling and the image) and $I(L;N)$ is maximised (there is optimum additional shared information with the other modality).

2.4 One Dimensional Example

These concepts can be illustrated with the following simple example. Consider the registration of two one-dimensional signals M and N shown in figure 5. The distance AB and $A'B'$ are equal and the dotted line on the lower figure shows the correct registration solution. The task is to register the two functions by maximising mutual information. In this simple example registration is by lateral translation only. Let the overall length corresponding to intensities m_1 and m_2 in M equal $x(m_1)$ and $x(m_2)$ respectively. Let the length (AB) corresponding to intensity n_2 in N equal $x(n_2)$. In addition we assume that the signal M is truncated as shown, that the signal N has sufficient extent to cover the whole of M for all possible registrations and that the registration solution is constrained so that AB must lie within the field of view of M.

The joint probability distribution is shown on figure 6 at registration. The intensity of each peak is proportional to the values of x indicated on the plot. This distribution and hence the measure of mutual information is identical to that which would occur if the feature AB in N were to overlay any position within $C'D'$ in M. The maximum of mutual

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31,
IEEE Computer Society Press, ISBN 0-8186-7367-2

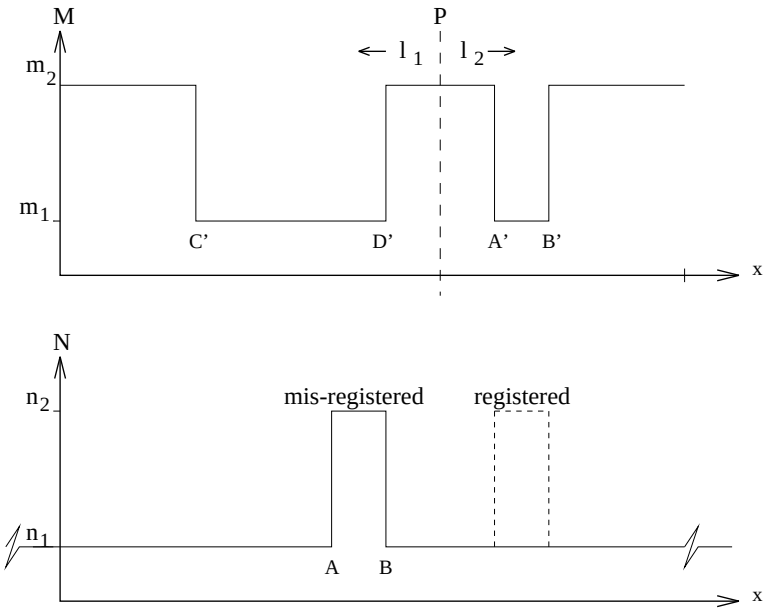

Figure 5. The registration of one dimensional signal $\mathbf{m}$ to $\mathbf{n}$ with additional labelling of M into $L\{l_1, l_2\}$

information is degenerate. In practice for an image with the addition of noise and other imaging artifacts there may well be an erroneous maximum of mutual information registering regions with similar intensities to those of the correct solution but corresponding to unrelated anatomical structures.

Now consider the addition of labels l_1 and l_2 as shown on figure 5. The partition, P, by the boundary between l_1 and l_2 is arbitrarily located on $D'A'$. The lengths of intensity m_1 and m_2 in region label l_1 are $x(m_1 l_1)$ and $x(m_2 l_1)$ respectively and in region label l_2, $x(m_1 l_2)$ $x(m_2 l_2)$ respectively. The joint probability distributions at registration are shown in the upper half of figure 7 and at misregistration with AB overlaying $C'D'$ in the lower half of figure 5. At registration there are only 4 peaks with intensities proportional to the lengths indicated while at misregistration there are 5. A few lines of algebra confirm that the mutual information of the correctly registered functions is now higher than the incorrectly registered functions.

This result holds for all positions of P in $D'A'$. We can therefore consider the partition into l_1 and l_2 as a coarse segmentation of the two regions.

3 Method

3.1 Image Data

A clinically acquired MR-PET image pair of the pelvis was used for the registration tests. Clinically the images were acquired for the localisation and staging of cervical cancer. The PET ^{18}FDG image of the pelvis was relatively low resolution, sampled with 57 $3.375mm$ slices of 128×128 $3mm \times 3mm$ pixels with a spatial resolution of around $8mm$ F.W.H.M. Regions of uptake usually take the form of small high contrast points in the image. ^{18}FDG also provides a weak differentiation between intra-abdominal tissues and fat and air but this is of low contrast and is influenced heavily by reconstruction artifacts and noise. Due to the different bed shapes in the the scanners this can also be non-rigidly deformed. The main source of anatomical information shared with MR is provided by $^{18}\mathrm{F}^-$ which is added as an additional tracer to identify bone structure. These bone landmarks are of greater contrast than the FDG tissue boundary but are still noticeably affected by reconstruction artifact and noise. Example axial and reformatted coronal slices from the PET image used in the test are shown in figure 8.

The MR data set was a clinically acquired T_1 weighted image consisting of 66 $4mm$ slices of 256×256 pixels of size $1.6mm \times 1.6mm$ as shown in figure 8. The MR image is of a much higher spatial resolution and contrast, showing many anatomical features. Many of these features, including marrow and fat, have a similar MR intensity with this sequence, as shown in figure 8.

3.2 Estimate of Registration

The image pair was registered by an experienced clinician interactively locating corresponding anatomical bone landmarks to provide a visually acceptable registration [9].

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31, IEEE Computer Society Press, ISBN 0-8186-7367-2

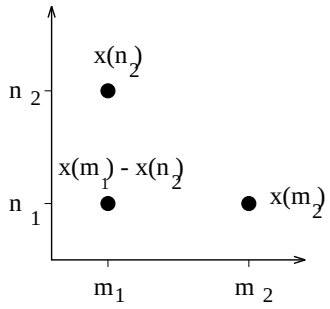

Figure 6. The two dimensional probability distribution of M and N at registration showing the relative intensities of 3 peaks.

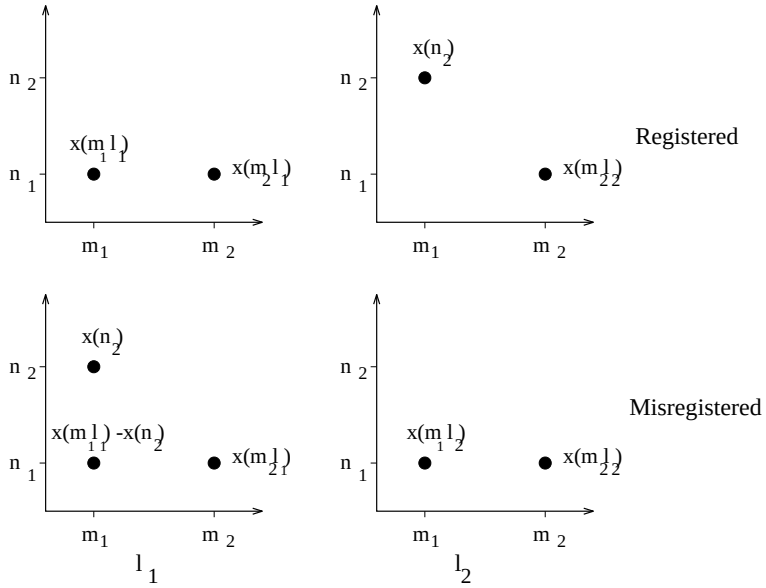

Figure 7. Slices through the three dimensional probability distribution of M (horizontal) and N (vertical) and labels (l_1 left and l_2 right) for registered (top) and unregistered (bottom) signals.

Due to the nature of the images the misregistrations between the modalities generally takes the form of large scale translational differences particularly along the length of the patient. Inspection of a range of manually estimated registration parameters for a number of datasets indicated typical misregistrations up to $90mm$ in translations along this axis.

Automated registration was attempted for the image pair shown in figure 8 by multi-resolution optimisation of mutual information $I(M;N)$ as described in [4]. This produced a poor estimate of the rigid registration parameters with a Z axis (along patient) translational error of greater than $30mm$.

3.3 Region Labelling

The MR image was labelled interactively into four categories:

- Air
- Fat
- Bone Marrow
- Non-Fatty Intra-abdominal tissue

This was carried out using an interactive intensity based region growing algorithm to produce a third image with voxels set to one of four values. Figure 9 illustrates a slice from the labelled volume.

3.4 Evaluation of $p\{m,l,n\}$

The MR image was first re-sampled at the same resolution as the PET image. Linear interpolation was used to increase axial sampling. A Gaussian kernel was used to blur

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31, IEEE Computer Society Press, ISBN 0-8186-7367-2

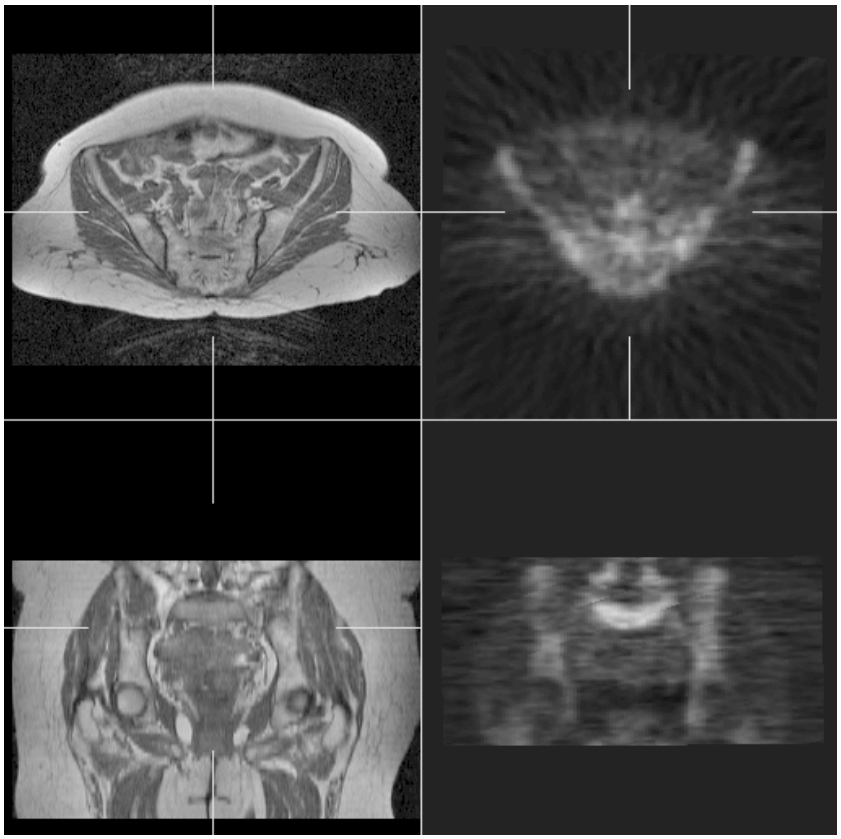

Figure 8. Example manually registered axial (top) and reformatted coronal (bottom) slices through an MR volume (left) and PET volume (right), illustrating the different representation of bone and fat in the two modalities.

the MR to give a similar spatial resolution to the PET. For a given rigid transformation T all voxels in the region of overlap of the two images were used to evaluate $p\{m, l, n\}$. The MR was used as the reference image (m) and corresponding PET (n) intensity was estimated using tri-linear intensity interpolation. MR and PET intensities were binned into 64 levels each and four bins were used for the labels in L to form a discrete distribution. Estimates of $I(M; N)$ and $I(M, L; N)$ were then evaluated from this distribution using equations 4 and 8 respectively.

4 Results

Figure 10 shows a plot of z axis displacement from the manual estimate of mutual information derived from only intensity, and of mutual information including labelling. The lower plot of $I(M; N)$ shows both a poor repose to misalignment and an optimum appreciably displaced from the manual estimate. The addition of labelling into the measure $I(M, L; N)$ increases the overall mutual information for all displacements produces distinct optimum close to the manual estimate.

A multi-resolution optimisation of $I(M, L; N)$ was carried out from a starting orientation defined by the scanner acquisitions, with the centres of the MR and PET volume aligned. The final transformation estimates are shown in table 1 along with those for optimisation of $I(M; N)$ and the manual estimate. The axes are as follows; x is from patient left to right, y is from patient front to back and z is from patient top to bottom. Registration by evaluation and optimisation of $I(M; N)$ derived from intensity only resulted in a visually poor registration. Multi-start of the optimisation and initialisation of the optimisation from the manual estimate confirmed that $I(M; N)$ was providing an incorrect global optimum corresponding to the alignment of bone features in PET with intra-abdominal fat. This resulted in both a large z axis displacement and a rotation around the y axis.

The manual estimate and the final estimate using the labelling approach $I(M, L; N)$ are within the expected accuracy of the manual registration procedure given the resolution of the original images. Interactive visual inspection of the results using colour overlay of PET intensity onto grey level MR in three orthogonal planes confirmed that the two results were visually comparable. Further work needs to be

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31,
IEEE Computer Society Press, ISBN 0-8186-7367-2

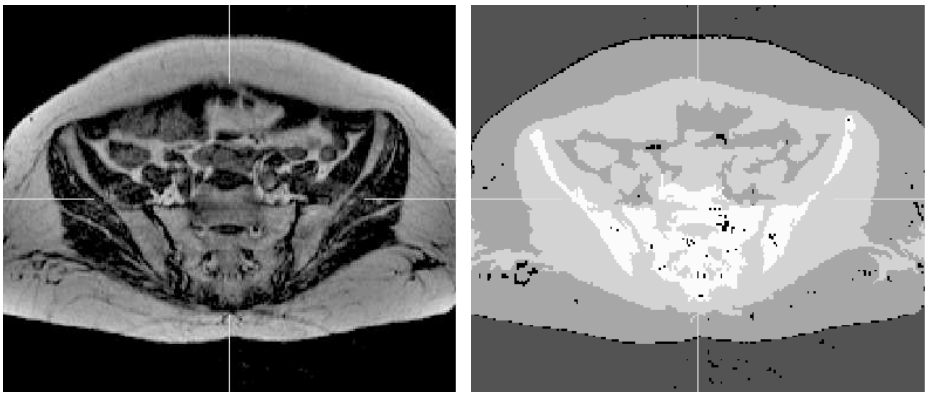

Figure 9. Example slice from MR volume (left) with labels (right).

done in assessing final accuracy.

Method	Translation (mm)			Rotation (deg.)		
	t_x	t_y	t_z	θ_x	θ_y	θ_z
Manual	-6.0	34.5	39.0	2.5	0.0	-2.0
$I(M;N)$	-4.3	34.3	15.6	2.4	-3.0	-0.2
$I(M,L;N)$	-4.9	34.1	37.1	3.9	0.0	0.0

Table 1. Rigid transformation parameters estimated for the test image pair

5 Discussion and Conclusion

In our test example we have interactively labelled voxel values to provide additional information in the registration measure. By using an information theoretic approach which includes these labels the segmentation need not be ideal allowing the possibility of using an automated labelling scheme. Segmentation schemes distinguishing strongly between unconnected regions are favoured. Work is underway in looking at how such schemes can be integrated into the registration process. In particular our current multi-resolution approach to optimising mutual information could be extended to include a multi-resolution region extraction stage.

Further work is required to assess the extent to which labelling can be corrupted before the the final accuracy (location of the global optimum) is affected. The effect of typical errors produced by automated labelling schemes on the registration measure for a range of clinical images needs to be determined.

In our simple 1D example and on the clinical test image we have supplied labelling information to differentiate regions of similar intensity which are distinguished in the other modality. Given a higher level knowledge about the regions it maybe be possible to use the same approach to combine regions of different intensity. This may be useful in reducing the influence of certain tissue boundaries which may have deformed between acquisitions.

This approach is an example of a case where pre-segmentation of images can provide additional information to help in the registration process. Alternatively it is also possible to use the information provided by the process of registration to assist in segmentation or classification.

We have introduced a new type of voxel similarity approach that makes use of voxel label information as well as the original image intensities. Coarse segmentation of the higher resolution image is necessary to provide the label information. Unlike a surface matching approach, eg: [2], this technique does not require identification of corresponding boundaries from both modalities, and we believe our information theoretic approach is more robust to segmentation errors. Van den Elsen et al [6] used an approach to MR and CT registration, in which an intensity mapping technique was used to make CT images "look" like MR images, prior to using cross correlation for registration. Intensity re-mapping in this way is only possible because of the simple relationship between tissue and intensity found in CT. The approach we propose here is more general, and preserves the original image intensities throughout. A possible interpretation of our technique is that the original intensity information is used for the accurate registration, and the labelling provides coarse scale information that removes unwanted optima far from the correct solution.

We have demonstrated this technique for the registration of MR and PET images of the pelvis, which could not be registered using the standard mutual information approach. The registration solution was close to that produced using interactively identified point landmarks and produced a visually acceptable result.

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31,
IEEE Computer Society Press, ISBN 0-8186-7367-2

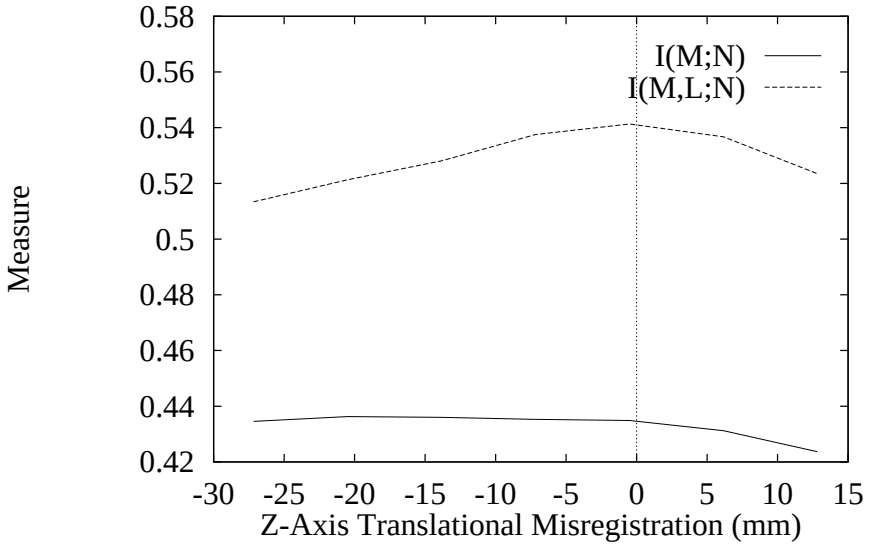

Figure 10. Graph of $I(M;N)$ and $I(M,L;N)$ with respect to axial (z) translation ($z = 0mm$ is the manual estimate).

6 Acknowledgements

This work was funded by the Guy's and St Thomas's NHS trust. We would particularly like to thank Prof. M. N. Maisey and the Clinical PET Centre of Guy's and St Thomas's. We are grateful for the support and encouragement of our clinical colleagues in this work, in particular Dr Joseph Wong (Radiologist), Dr Wai-Lup Wong (Radiologist), and for the technical assistance of the Radiographic staff of Guy's and St. Thomas' Hospitals in London. We would also like to thank Dr John Little of UMDS for mathematical advice and we acknowledge useful discussions with Andre Collignon and Dirk Vandermeulen of KUL, Leuven, Belgium and Sandy Wells and Paul Viola of MIT and Brigham and Women's Hospital Boston, USA, on the subject of mutual information.

References

[1] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multimodality image registration using information theory. In B. Y., B. C, and D. P. R., editors, *Proceedings of Information Processing in Medical Imaging*, pages 263–274, 1995. Brest, France.

[2] H. Jiang, R. Robb, and K. Holton. New approach to 3-D registration of multimodality medical images by surface matching. In *Proceedings of Visualisation in Biomedical Computing*, pages 196–213, 1992. Chapel Hill, N.C.,U.S.A.

[3] F. Reza. *An Introduction to Information Theory*. Dover, New York, 1994.

[4] C. Studholme, D. Hill, and D. Hawkes. Automated 3D registration of truncated MR and CT images of the head. In D. Pycock, editor, *Proceedings of the British Machine Vision Conference*, pages 27–36. BMVA, 1995. Birmingham.

[5] C. Studholme, D. Hill, and D. Hawkes. Multiresolution voxel similarity measures for MR-PET registration. In B. Y., B. C., and D. P. R., editors, *Proceedings of Information Processing in Medical Imaging*, pages 287–298, 1995. Brest, France.

[6] P. Van den Elsen, E. Pol, T. Sumanawaeera, P. Hemler, S. Napel, and J. Adler. Grey value correlation techniques used for automatic matching of CT and MR brain and spine images. In *Proceedings of Visualisation in Biomedical Computing*, pages 227–237, 1994. Rochester Mn.,U.S.A.

[7] P. Viola and W. Wells. Alignment by maximisation of mutual information. In *Proceedings of the 5th International Conference on Computer Vision*, pages 15–23, 1995.

[8] W. Wells, P. Viola, and R. Kikinis. Multimodal volume registration by maximization of mutual information. In *Proceedings of the 2nd annual international symposium on Medical Robotics and Computer Assisted Surgery*, pages 55–62, 1995. Baltimore, U.S.A.

[9] W. Wong, C. Studholme, P. Lewis, K. Raju, R. Beaney, K. Tonge, T. Nunan, D. Hawkes, and J. Pemberton. Combined MR,CT and PET imaging in oncological patients. *British Journal of Radiology*, 66(suppl):33–34, 1993.

Published in: Proceedings of the IEEE workshop on Mathematical Methods in Biomedical Image Analysis,San Francsico Ca., June 1996, pp. 23-31, IEEE Computer Society Press, ISBN 0-8186-7367-2