

Multi-channel Mutual Information Using Scale Space

Mark Holden¹, Lewis D. Griffin¹, Nadeem Saeed², and Derek L.G. Hill¹

¹ CISG, Imaging Sciences, GKT School of Medicine, Guy's Hospital, King's College
London, UK

`mark.holden@kcl.ac.uk`

² GlaxoSmithKline, MRI Unit, The Frythe, Welwyn, UK

Abstract. We propose a new voxel similarity measure which utilises local image structure and intensity information. Gaussian scale space derivatives provide the structural information. Each derivative is assigned an information channel of N-D normalised mutual information. We illustrate the behaviour of the measure for a simulated signal and 2D medical brain images and demonstrate its potential for non-rigid, inter-subject registration of 3D thorax MR images as a proof of concept.

1 Introduction

Voxel intensity based image similarity measures, particularly mutual information, perform well for intramodality [1] and intermodality [2] rigid-body registration. However, non-rigid registration is a more difficult problem because the deformations that need to be recovered can be localised and have high frequency content which requires a high dimensional spatial transformation. As the dimensionality of the transformation increases so does the the likelihood of false optima of the similarity measure. Transformations with local support provide a useful model of tissue deformation, but as the dimensionality increases the number of voxels in the support region decreases which leads to a reduction in local information available for measuring similarity. False or local optima lead to ill-conditioning of the registration function which increases the possibility of a registration algorithm getting trapped in a local optima. There are approaches to reduce the likelihood of this by including additional constraints derived from modelling the physical tissue deformation or by regularising the transformation to reduce its dimensionality. The physical modelling approach has the disadvantage of requiring a labelled tissue model and the mechanical properties of the tissue which could vary between individuals and over time, for instance during disease progression or regression. Regularisation can restrict the solution space of transformations to satisfy certain mathematical constraints. This is definitely advantageous so long as the transformation adequately models the physical deformation. However, it cannot prevent the algorithm getting trapped in local optima. Here we investigate an approach aimed at improving the similarity measure by incorporating additional local structural image information.

Related Work

Shen et al. [3] have designed a similarity measure that determines image similarity based on a attribute vector for each voxel at GM, WM and CSF interfaces. The attribute vector is derived from the voxel's edge type and geometric moment invariants calculated from voxel intensities in a spherical neighbourhood. This similarity measure is specifically designed for intra-modal, inter-subject MR brain image registration and requires a GM, WM and CSF segmentation.

In contrast, we aim for a general purpose registration algorithm that can be applied to intermodality data direct from the scanner without a pre-processing step. We start by establishing a set of desirable properties of the similarity measure and use these to devise a mutual information measure that utilises more structural image information than simple intensities. In this way we can retain the desirable intermodality property of mutual information. We use the derivatives of the Gaussian scale space expansion of the image to provide this local information. To assess the performance of the measure we present some simulations and results of inter-subject intramodality registration experiments.

2 Problem Analysis

Properties of Non-rigid Registration Similarity Measures

We require metrics that: (a) are suitable for intermodality data; (b) are translation and rotation invariant (invariance to non-rigid motion might also be desirable); (c) return values that are a smooth decreasing function of misregistration.

A problem often encountered in non-rigid registration arises when non-corresponding parts of the anatomy have similar MR signal intensities. When these regions overlap, voxel intensity based similarity measures can give a local optima (c.f. non-corresponding gyri [3]. Some authors have address this using additional information Cachier et al. [4]. To avoid this additional spatial information is required. Shen et al. [3] used geometrically invariant moments, Plum [5] used intensity and first derivative, Butz [6] used a feature space. We want to use local image structure to provide this information. Ideally, we would like a small set of measures that unambiguously characterise each image voxel. To facilitate optimisation we desire measures that can be used in multi-resolution search, c.f. rigid-body registration [2].

Incorporating Local Image Structure in Mutual Information

Consider the Taylor series expansion of a 3D function $f(x, y, z)$ about (a, b, c) :

$$f(x, y, z) = \sum_{l=0}^{\infty} \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \left\{ \frac{\partial^{lmn} f(x, y, z)}{\partial x^l \partial x^m \partial x^n} \right\}_{a,b,c} \frac{(x-a)^l}{l!} \frac{(y-b)^m}{m!} \frac{(z-c)^n}{n!} \quad (1)$$

This equation describes how a continuous function $f(x, y, z)$ can be expanded in terms of its derivatives at (a, b, c) . Since we are interested in discrete images we instead propose to use the analogous Gaussian scale space expansion [7].

$$I(x, y, z) = I_0(x, y, z, \sigma) + \sum_{i=x,y,z,\sigma} C_i I_i + \sum_{i,j=x,y,z,\sigma} C_{ij} I_{ij} + \dots \quad (2)$$

Where I_0 denotes the image convolved with a Gaussian of standard deviation σ and $I_{ij}(x, y, z, \sigma)$ denotes the derivative of I_0 w.r.t. to the three Cartesian directions i, j and scale σ . It has been shown that such a representation maximally captures local image structure. The derivative terms are almost statistically independent. The small amount of dependence is not a problem.

We need a multivariate similarity measure that uses the derivative terms of the scale space expansion of the image. One way to do this is to consider each derivative term (I_{ij}) as a separate information channel. Then we can use multi-dimensional mutual information on the joint event $\{I_{ij}\}$. For two information channels, we get a 4D joint histogram and mutual information of the form:

$$MI(A, B, C, D) = H(A, B) + H(C, D) - H(A, B, C, D) \quad (3)$$

$$NMI(A, B, C, D) = \frac{H(A, B) + H(C, D)}{H(A, B, C, D)} \quad (4)$$

Where A, B and C, D refer to derivatives of the target and source respectively.

3 Materials and Methods

3.1 Implementation

Gaussian scale-space: In our experiments we consider only the luminance, first and second order derivative terms of the scale space expansion. These are reasonably invariant to rigid motion. Invariance to non-rigid motion could be achieved by recomputing them each time similarity is measured, however this would increase the computational overhead. The luminance image $I_0(\mathbf{x})$ is generated by convolving the image $I(\mathbf{x})$ with a Gaussian kernel $G(\mathbf{x})$ viz: $I_0(\mathbf{x}) = G(\mathbf{x}) \odot I(\mathbf{x})$ where $G(\mathbf{x}) = \frac{1}{2\pi\sigma^2} \exp(-|\mathbf{x}|^2/2\sigma^2)$. The gradient magnitude image $I_1(\mathbf{x}) = |\nabla(I_0)|$ and the Laplacian image $I_2(\mathbf{x}) = \nabla^2(I_0)$. In the experimental work we refer to these as luminance, GLOG (gradient magnitude of luminance) and LOG (Laplacian of luminance). The intensity of the LOG image was normalised by subtracting the minimum so that its minimum is zero. To avoid truncation during convolution, the image was reflected about each boundary by half the kernel width. Gaussian convolution and differentiation (central derivative) were implemented in matlab (Mathworks Inc, MA, USA) for 1D signals and 2D images and in C++ using vtk (Kitware, NY, USA) classes for 3D images. In all instances, the kernel radius was chosen to be three times larger than the standard deviation to avoid truncation effects.

Multi-dimensional mutual information: A major difficulty with this is that the dimensionality of the joint histogram array depends on the number of derivative terms (n). The array size grows as a power of n . This can lead to a sparsely populated array, also the memory required and access time grow as a

power of n . Reducing the number of bins can help, but this only results in a linear reduction of size.

Image interpolation is generally the most computationally intensive part of voxel-based algorithms and grows linearly with n . A possible way of reducing the overhead could be to down-sample images. For 3D images, down-sampling by a factor of 2 reduces the number of voxels that need to be interpolated by a factor $2^3 = 8$. In summary, this approach seems viable for small n with down-sampling.

All similarity measures were implemented in both matlab (1D and 2D) and also in C++ for 3D images. For the non-rigid registration of 3D images a segmentation propagation algorithm based on method described in [8] and the 4D similarity measures were implemented in C++ and vtk by redesigning a number of classes of the CISG registration toolbox [9].

3.2 Simulation: Geometric Scaling of Synthetic Signal

A test signal was constructed by low-pass filtering a signal consisting of two rectangular pulses. We chose to model the imaging system using a unit width Gaussian low pass filter. The luminance, gradient magnitude of luminance GMOG and Laplacian of luminance LOG signals were generated from the test signal using a Gaussian filter of standard deviation $\sigma = 6$ samples. To assess the behaviour of similarity measures as a function of misregistration (registration function) a copy of the test signal was geometrically scaled relative to the original signal. The similarity of these two signals was measured as a function of the scale factor ($s_x, 1 \leq s_x \leq 3$), where $s_x = 1$ represents perfect registration.

Figure 2 shows the resulting graph for four similarity measures: standard normalised mutual information (NMI), NMI applied to luminance signal, 4D NMI using luminance and GMOG, 4D NMI using luminance and LOG.

For the standard form, there is a false maximum at $s_x \approx 1.6$ and the function is ill-conditioned for $s_x > 1.6$. Gaussian smoothing helps condition the registration function, but the function is flat around $s_x = 1.9$. For the 4D measures, both are well-conditioned and relatively easy to optimise.

3.3 Experiment: Translational Misregistration of a Brain Sub-image

This experiment was designed to evaluate the behaviour of our proposed similarity measure by taking two 2D images of the same anatomy and misregistering a small sub-image of one relative to the other. The data was acquired by scanning a volunteer's brain with a special T1W 3D gradient echo MR sequence with two interleaved readout lines. This data was reconstructed into two 3D spatial images separated by an interval of TR (a few milli-seconds). Essentially the difference between the two images is noise, but there is also a small difference in motion artefacts due to fast flowing blood. These images can be considered as a registration gold-standard, and the graphs of the registration function tell us how the similarity measure behaves as a function of misregistration for images with a noise difference. We took an axial slice through the lateral ventricles and




Fig. 1. (a) Test signal used for registration simulation experiments. (b) luminance signal, filtered with a Gaussian ($\sigma = 4$) and gradient magnitude of luminance (GLOG) ($\sigma = 4$). (c) luminance and Laplacian of luminance (LOG) ($\sigma = 4$).




Fig. 2. Normalised mutual information (NMI) as a function of geometric scale change ($s_x, 1 \leq s_x \leq 3$). Comparison of standard form; standard form with Gaussian blurring $\sigma = 6$; 4D NMI with Gaussian and gradient magnitude of luminance (GLOG) input channels ($\sigma = 6$); 4D NMI with Gaussian and Laplacian of luminance (LOG) input channels ($\sigma = 6$).

extracted a 32×32 pixel sub-image. Then we misregistered the sub image relative to the other image by applying uniform scaling using a scale factor s_x in the range $1 \leq s_x \leq 3$ where $s_x = 1$ represents perfect registration.

Figure 3 shows the results of the experiment. The standard NMI flattens out for $s_x > 2$ making it difficult to optimise. For $\sigma = 1$ voxel (Figure 3 (a)), there is little difference between the other measures performance. However, for $\sigma = 8$ voxels there is a large difference. Standard 2D NMI applied to the luminance image has an optima around ($s_x = 1.8$) and is flat for ($s_x > 2.5$) whereas the 4D NMI measure based on the luminance and LOG has the least ill-conditioning and widest capture range. This behaviour could be important for multi-resolution optimisation, thought necessary for recovering large deformation.

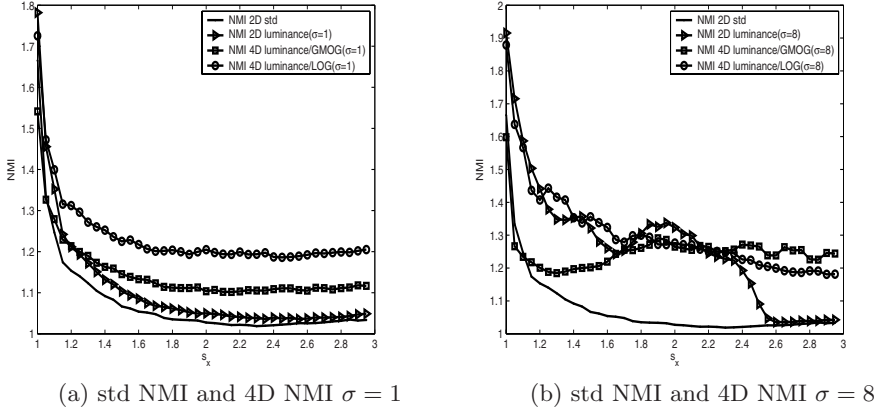


Fig. 3. Plots of the similarity as a function of geometrical scale misregistration of four similarity measures using different amounts of Gaussian blurring. (a) standard NMI (no blurring); standard NMI for Gaussian blurred (luminance) images ($\sigma = 1$); 4D NMI luminance and GMOG ($\sigma = 1$); luminance and LOG ($\sigma = 1$). (b) standard NMI; standard NMI applied to luminance images ($\sigma = 8$); 4D NMI luminance and GMOG ($\sigma = 8$); luminance and LOG ($\sigma = 8$).

3.4 Experiment: Non-rigid Intersubject Registration of the MR Lung

This experiment was designed to evaluate the accuracy of the similarity measures for a realistic 3D non-rigid registration problem involving large deformations. We used T1W 3D gradient echo MR scans of the thorax of seven murine subjects¹. These images contained 128^3 isotropic voxels of dimension 0.234mm. The chest wall was manually segmented in each image by an expert with knowledge of the anatomy. We assigned one of these images as the atlas and registered this to the six other subjects using the segmentation propagation algorithm. We used the transformations to propagate the atlas segmentation into the space of the six other subjects. Because of the large difference in repositioning and subject size registration was initialised with a manually determined scaling transformation: six rigid-body plus three scale parameters. The registration strategy was based on a rigid-body stage followed by three B-Spline non-rigid stages with control point spacing of 5, 2.5 and 1.25 mm. To evaluate the accuracy of segmentation propagation we compared the propagated lungs with the manually defined segmentation using the overlap similarity index $S = \frac{2N(R1 \cap R2)}{N(R1) + N(R2)}$ as used in [8]. Where $R1$ and $R2$ are the sets of voxels in the propagated and segmented regions and $N(\cdot)$ refers to the number of voxels in a set. Table 1 shows similar performance of the 4D measure compared to the standard 2D one and indicates a similar accuracy for down-sampled and original resolution images.

¹ Animals were housed, maintained and experiments conducted, in accordance with the Home Office Animals (Scientific Procedures) Act 1986, UK.

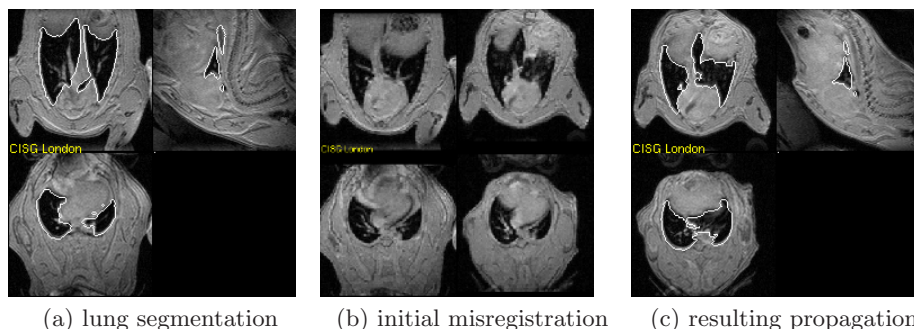


Fig. 4. (a) Manual segmentation of lungs. (b) Example of initial misregistration after manual alignment. (c) Example of a segmentation propagation.

Table 1. Overlap similarity index S after manual alignment (manual) and after rigid-body and non-rigid registration with B-spline control point spacing of 5mm, 2.5mm and 1.25mm. This table compares standard normalised mutual information and the 4D form using luminance and gradient magnitude of luminance GMOG and luminance and Laplacian of luminance (GLOG). The luminance image is created by filtering with Gaussian of $\sigma = 2$ and $\sigma = 4$. The three left hand columns refer to original resolution data and the three right hand ones to images down-sampled by a factor of 2.

target	manual	original resolution			down-sampled		
		NMI2D	GLOG	GMOG	2d	GMOG	GL
		$\sigma = 0$	$\sigma = 2$	$\sigma = 4$	$\sigma = 4$		
subject 1	0.36	0.82	0.76	0.75	0.79	0.72	0.73
subject 2	0.47	0.87	0.80	0.78	0.83	0.71	0.73
subject 3	0.41	0.87	0.81	0.80	0.83	0.74	0.76
subject 4	0.53	0.89	0.88	0.88	0.85	0.72	0.79
subject 5	0.26	0.87	0.86	0.86	0.84	0.78	0.81
subject 6	0.45	0.84	0.81	0.78	0.81	0.76	0.76

4 Discussion and Conclusions

We have established a set of properties of similarity measures for non-rigid inter-modality image registration. We used these to design a novel similarity measure based on Gaussian scale space derivatives. We demonstrated that this has a wider capture range than standard forms for large deformations with synthetically misregistered signals. We showed satisfactory performance for translating a sub-image of 2D brain slices. For non-rigid registration of 3D lung data there is similar performance with the standard measure for inter-subject registration accuracy, as assessed by comparing the overlap of propagations with manual segmentation. The lung images had little contrast between different tissues. This may be a reason why the 4D measure did not perform better than the standard one.

Acknowledgements. We are grateful to GlaxoSmithKline for funding this work.

References

1. M. Holden et al. Voxel similarity measures for 3D serial MR brain image registration. *IEEE Transactions on Medical Imaging*, 19(2):94–102, 2000.
2. C. Studholme, D. L. G. Hill, D. J. Hawkes. Automated 3-D registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Medical Physics*, 24(1):25–35, 1997.
3. Dinggang Shen; Davatzikos, C.; Dinggang Shen; Davatzikos, C. HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Transactions on Medical Imaging*, 21(11):1421–1439, November 2002.
4. P. Cachier et al. Multisubject Non-Rigid Registration of Brain MRI using Intensity and Geometric Features. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI '01)*, pages 734–742, Utrecht, Netherlands, 2001.
5. J. P. W. Pluim et al. Image registration by maximization of combined mutual information and gradient information. *IEEE TMI*, 19(8):809–814, 2000.
6. T. Butz et al. Affine registration with Feature Space Mutual Information. In *MICCAI '01*, pages 549–556, Utrecht, Netherlands, 2001.
7. L.M.J. Florack et al. Cartesian Differential Invariants in Scale-Space. *J. Math. Imaging Vision*, 3(4):327–348, 1993.
8. M. Holden, J. A. Schnabel, D. L. G. Hill. Quantification of small cerebral ventricular volume changes in treated growth hormone patients using non-rigid registration. *IEEE Transactions on Medical Imaging*, 21(10):1292–1301, 2002.
9. T. Hartkens et al. *VTK CISG Registration Toolkit*. CISG, Imaging Sciences, King's College London. <http://www.imageregistration.com/>.