

Comparison of Popular Non-Rigid Image Registration Techniques and a New Hybrid Mutual Information-Based Fluid Algorithm

C. Fookes and A. Maeder
Research Concentration for Computer Vision and Automation
Queensland University of Technology,
GPO Box 2434 Brisbane, 4001 QLD Australia
c.fookes@qut.edu.au

Abstract

Recently there has emerged a need to compute multimodal non-rigid registrations in a lot of clinical applications. To date, the viscous fluid algorithm is perhaps the most adept method at recovering large local mis-registrations that exist between two images. However, this model can only be used on images from the same modality as it assumes similar intensity values between images. This paper presents a solution to this problem by proposing a hybrid non-rigid registration using the viscous fluid algorithm and mutual information (MI). The MI is incorporated via the use of a block matching procedure to generate a sparse deformation field which drives the viscous fluid algorithm. This algorithm is compared to two other popular local registration approaches, namely Gaussian convolution and the thin-plate spline warp. Results show that the thin-plate spline warp and the MI-Fluid approach produce comparable results. However, Gaussian convolution is the superior choice, especially in controlled environments.

1. Introduction

Non-rigid image registration is an essential tool required for overcoming the inherent local anatomical variations that exist between images acquired from different individuals or atlases. The majority of these non-rigid algorithms assume the existence of similar intensities between images, restricting their use to intra- or mono-modality registrations. Recently, however, there has emerged a need to compute multimodal non-rigid registrations in a lot of clinical applications. The most prominent application of this is in the registration of pre-operative and intra-operative images. This allows the display of pre-operative anatomical and pathological tissue discrimination in the interventional field [7].

An important concept that arises in the computer vision

field during the mid 1990's was an entropy-based measure known as mutual information (MI). This measure has its roots in information theory and has demonstrated its power and robustness for use in multimodality registration in the rigid domain repeatedly. The strength of this measure lies in its simplicity as it does not assume the existence of any particular relationship between image intensities. It only assumes a statistical dependence.

MI has been incorporated into a non-rigid registration by several researchers. The main distinction between the proposed methods lie in the way the MI is calculated. This is accomplished either globally or locally [4]. However, to date MI has never been incorporated with a physical continuum model, (such as the elastic or viscous fluid algorithm). The viscous fluid algorithm is a popular approach which is capable of recovering large local mis-registrations. It also ensures that the deformation field is physically smooth. However, like most other non-rigid registrations, it assumes similar intensity values between images.

This paper proposes a novel hybrid non-rigid registration using the viscous fluid algorithm and MI. This new technique is also compared to two other popular non-rigid registration approaches, namely Gaussian convolution and the thin-plate spline warp. All three methods rely on the execution of a block matching procedure to generate an initial sparse deformation field. However, the way in which this sparse deformation field is propagated to the rest of the image depends on the technique utilised.

The outline of the paper is as follows. Some MI preliminaries are outlined in Section 2. Section 3 introduces non-rigid image registration in general, while Section 4 describes the techniques examined by this paper. This includes a general block matching approach, Gaussian convolution, thin-plate spline warps, and the new hybrid algorithm incorporating MI and the viscous fluid algorithm. Results are presented in Section 5 and conclusions are drawn in Section 6.

2. Mutual Information Preliminaries

MI is an information theoretic measure and was proposed for use in image registration by two independent groups, Viola et al. [10] and Collignon et al. [3], in 1995. The basic concept behind the use of this measure is to find a transformation, which when applied to an image, will maximise the MI between the two images.

The MI formulation used by the techniques described in this paper is based on the Kullback-Leibler measure [9] and is given by,

$$I(X, Y) = \sum_{x, y} p_{X, Y}(x, y) \log \left(\frac{p_{X, Y}(x, y)}{p_X(x) p_Y(y)} \right) \quad (1)$$

where the densities are estimated by normalisation of the 2D frequency histograms. MI is a measure of the degree of dependence of the random variables X and Y . When formulated using the Kullback-Leibler measure in Equation 1, the MI measures the distance between the joint distribution $p_{X, Y}(x, y)$ and the distribution associated with complete independence, i.e. $p_X(x) \cdot p_Y(y)$ [8]. This measure is bounded below by complete independence and bounded above by one-to-one mappings.

3. Non-Rigid Image Registration

A rigid registration is composed solely of a rotation and translation and literally preserves the ‘rigid’ body constraint, i.e. a body is rigid and must not undergo any local variations during the transformation. This type of registration is distance preserving and is adequate for many applications in medical imaging including multimodality and intra-patient registration. However, for inter-patient registration or patient-atlas matching, non-rigid algorithms are required. In a non-rigid approach, the ‘rigid’ body constraint is no longer acceptable as it does not account for the non-linear morphometric variability between subjects [6], i.e. there exists inherent anatomical variations between different individuals resulting in brain structures that vary in both size and shape. These non-rigid algorithms allow one image to deform to match another image, thus overcoming any local variations.

A non-rigid registration defines a deformation field that gives a translation or mapping for every pixel in the image. This is generally described by the following relationship.

$$I_f \circ T(\mathbf{x}) = I_f(\mathbf{x} - \mathbf{u}(\mathbf{x})) = I_r \quad (2)$$

In the above expression, I_f is referred to as the floating image that is undergoing the deformation while I_r is the reference image. T denotes the non-rigid transformation which equates to a translation of every pixel \mathbf{x} in the floating image by a certain displacement defined by the displacement field $\mathbf{u}(\mathbf{x})$.

4. Description of Techniques

There are many ways of estimating the required displacement field $\mathbf{u}(\mathbf{x})$ in Equation 2. This includes deformable models, optical flow, elastic and viscous fluid models, spline warps, truncated basis function expansion methods, and also local registration approaches [4]. The type of method employed will also determine what constraints are imposed on the deformation field. Generally speaking, the constraints are used to ensure the existence of a smooth and continuous deformation field.

The techniques that will be described here however, are all based on a local registration approach referred to as block matching. This method is quite popular as it easily allows the incorporation of the MI measure into the non-rigid registration. This approach is described below, along with the three techniques which are used to propagate the sparse deformation field to the entire image. They are Gaussian convolution, the thin-plate spline warp, and a new hybrid algorithm incorporating MI and the viscous fluid algorithm.

4.1. Block Matching

Non-rigid registration can be made possible through local registration approaches and several methods exist to accomplish this. One common method, known as block matching, is where a grid of control points are defined on an image which are each taken as the centre of a small window. These windows, which usually overlap their neighbours, are then translated to maximise a local similarity criterion. MI is used as the similarity measure in order to obtain a robust multimodality non-rigid registration.

The location of the maximum can then be found through an exhaustive search or with the use of local optimisation strategies. The location of the maximum then represents the existence of a corresponding window in the second image, the centre of which being the homologue point of the corresponding grid point defined in the first image. Thus, this block matching approach can be used to generate two corresponding sets of control points (or landmark points) between two images. This information can then be used to generate a sparse deformation field with the translations known at each of these grid points. An example of a sparse field generated using block matching procedures is shown in Figure 1.

4.2. Gaussian Convolution

As described above, the execution of a block matching procedure results in the generation of two corresponding sets of control points. By using these control points with known deformations in a non-rigid registration, constraints are being imposed on the space of possible deformations.




Figure 1. Sparse deformation field calculated using a block matching procedure.

This has been described as a static constraint problem [5], or an interpolation issue as the problem then becomes one of how to interpolate the deformations at these known locations to the rest of the image. Several techniques exist to accomplish this.

One of the simplest approaches is to convolve this sparse deformation field with a 2D Gaussian kernel (Gaussian smoothing), to propagate the deformations to the rest of the image. It has been described in [6] that Gaussian smoothing is equivalent to solving a heat or diffusion equation. Thus, this approach equates to an oversimplified version of a physical model-based algorithm (such as the elastic or viscous-fluid model). As model-based techniques are solved in an iterative process, the two choices essentially become whether to perform Gaussian smoothing on either the final or incremental deformation field. The first choice equates to an oversimplified elastic transformation while the second choice equates to an oversimplified viscous fluid transformation [6].

4.3. Thin-Plate Spline Warp

Another popular approach very suited to the propagation of a sparse deformation field is the thin-plate spline warp. In this method, an image is represented as a thin metal plate which undergoes certain deformations at selected points, defined by the sparse deformation field. The thin-plate spline has an elegant algebra that expresses the dependence of the physical bending energy of the thin metal plate to these point constraints [1].

For 2D image registration, two 2D thin-plate spline warps are used to describe an interpolation map from R^2 to R^2 relating two sets of landmark points, (one for the deformation in the x and y -directions respectively). The fundamental basis function used by the thin-plate spline is given by the following expression,

$$z(x, y) = -U = -r^2 \log r^2 \quad (3)$$

where r is the distance $\sqrt{(x^2 + y^2)}$ from the Cartesian ori-

gin. The function $U(r)$ also satisfies the following equation.

$$\delta^2 U = \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right)^2 U \propto \delta_{0,0} \quad (4)$$

Thus, U is a fundamental solution of the biharmonic equation $\Delta^2 U = 0$, the equation for the shape of a thin metal plate vertically displaced as a function $z(x, y)$ above the (x, y) -plane. Note that this basis function is the natural generalisation to two dimensions of the function $|x|^3$ which describes the common 1D cubic spline [1].

A thin metal plate which is subjected to vertical displacements at selected points with any arbitrary spacing will minimise the 2D bending energy of the metal plate. This is equivalent to minimising the following expression.

$$\int \int_{R^2} \left(\left(\frac{\partial^2 z}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 z}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 z}{\partial y^2} \right)^2 \right) dx dy \quad (5)$$

The minimisation of this energy represents a smoothness criterion which imposes constraints on the deformation field, ensuring that the deformation in between the known landmark points varies smoothly. Note that this process is repeated twice - for the deformation in the x and y directions respectively.

4.4. A New Hybrid MI-Based Fluid Algorithm

To date, the viscous fluid registration algorithm is perhaps the most adept method at recovering large local misregistrations that exist between two images. This is due to the internal restoring forces which relax as the image deforms over time. This method ensures that the deformation field is physically smooth. However, like the elastic model, the viscous fluid model can only be used on images from the same modality as it assumes similar intensity values between images.

In the viscous fluid model, the instantaneous velocity field $\mathbf{v}(\mathbf{x}, t)$ is linked to external forces by the Navier-Stokes viscous fluid partial differential equation which is shown below [2],

$$\alpha \nabla^2 \mathbf{v}(\mathbf{x}, t) + (\alpha + \beta) \nabla (\nabla^T \cdot \mathbf{v}(\mathbf{x}, t)) + \mathbf{b}(\mathbf{x}, \mathbf{u}(\mathbf{x}, t)) = 0 \quad (6)$$

where $\mathbf{v}(\mathbf{x}, t)$ is the instantaneous velocity of the displacement field $\mathbf{u}(\mathbf{x}, t)$ at time t . The term $\mathbf{b}(\mathbf{x}, \mathbf{u}(\mathbf{x}, t))$ represents the applied forces and the parameters α and β are the viscous fluid coefficients. This equation is solved at each time step and the driving forces are derived from image differences and intensity gradients.

The main motivation behind the creation of a hybrid algorithm was to incorporate the strengths of both the viscous

fluid algorithm and an information theoretic measure such as MI. This would allow the execution of a fluid registration on multimodal images. In the original viscous fluid algorithm described above, the driving forces are formulated in the most possible local manner, i.e. the force acting at a particular voxel is derived from the intensity difference and gradients of a point, not a region. However, in the approach of the hybrid algorithm, these driving forces are replaced with those derived from the MI block matching scheme. As mentioned in Section 4.1, the block matching is used to produce two sets of corresponding point sets with known deformations at each point. MI is the similarity criterion used in order to allow for a multimodal registration. The MI is also formulated using the frequency histogram approach and Equation 1.

The forces derived from the sparse deformation field are then fed into the viscous fluid algorithm which are used as the driving potentials instead of the original image differences and gradients. The significant difference however, lies in the manner in which the forces were calculated. Instead of utilising the intensity difference and gradients of a point, the block matching approach estimates the displacement field and hence the driving forces of a point by incorporating information that is contained in a small region around the point.

5. Results

The three local registration approaches were tested on a pair of simulated multimodal images with known deformations. The simulated multimodal images were generated from a single MR image and a deformed version of itself. The intensities of the reference image were also transformed using $I^* = \sin(I \times \frac{\pi}{255})$ to simulate images from different imaging modalities. These images are shown in Figure 2, along with a rescaled difference image. This difference image however, was computed without the intensity transformation in order to display more meaningful results. This is also the case for other difference images displayed later.

The results of the registration are shown in Figure 3. The letters (a), (b), and (c) are used to represent results computed with Gaussian convolution, the thin-plate spline warp, and the MI-Fluid algorithm respectively. The numbers (1), (2), and (3) are used to represent the final image after registration, the rescaled difference image, and a histogram of the intensity differences respectively. Quantitative results are shown in Table 1. This includes the SSD (sum of square differences) and SAD (sum of absolute differences) measures, and the mean μ and standard deviation σ of the error. These results are also shown for the two images before registration described by the term ‘pre-reg’.

From the rescaled difference images shown in Figure

3, it appears that all three algorithms have reduced local anatomical differences quite considerably when compared to the differences before registration. However, these rescaled difference images can be a little misleading as the intensities representing the differences are scaled to fit into the range $\{0 - 255\}$, no matter how large the actual difference. Thus, the histogram of intensity differences is also presented as another helpful avenue for evaluating the results.

Method	SSD	SAD	Error μ	Error σ
Pre-Reg	1.21×10^7	2.19×10^5	0.389	185.21
Gauss Conv	4.16×10^4	1.17×10^4	-0.003	0.64
TPS	2.52×10^6	9.43×10^4	0.099	38.48
MI-Fluid	2.82×10^6	6.49×10^4	0.109	42.98

Table 1. Quantitative error measures of registration results.

From the histograms, it can be seen that all methods have errors concentrated around the origin. However, Gaussian convolution has a much lesser spread of its errors than the other two approaches. This is also illustrated in Table 1. The Gaussian convolution method also has significantly lower SSD and SAD scores, as well as a mean error closer to the origin, and a much smaller error standard deviation. The MI-Fluid algorithm has a larger SSD score than the thin-plate spline warp, yet it has a smaller SAD score. This suggests that overall, the MI-Fluid approach produces less errors than the thin-plate spline warp. However, MI-Fluid has more errors in the outer regions which carry more weighting in the SSD measure, resulting in a higher SSD score than the thin-plate spline warp. Other results between these two methods are comparable.

6. Conclusion

This paper has proposed a hybrid non-rigid registration algorithm using MI and the viscous fluid algorithm. The MI is incorporated via the use of a block matching procedure to generate a sparse deformation field which drives the viscous fluid algorithm. Results show that the hybrid approach is successful in recovering local deformations between multimodal images. However, it is susceptible to interpolation artifacts which prevent the estimation of sub-pixel translations. Thus, the estimated deformation field will not vary smoothly, instead it will vary with integer valued steps.

This algorithm was also compared to two other popular local registration approaches, namely Gaussian convolution and the thin-plate spline warp. Results showed that the thin-plate spline warp and the MI-Fluid approach produced comparable results. However overall, simple Gaussian convolution was significantly superior. The main drawback of using Gaussian convolution is that appropriately sized variances




Figure 2. Manually deformed simulated multimodal images. (a) Reference image with intensity transformation $I^* = \sin(I \times \frac{\pi}{255})$, (b) Deformed floating Image, (c) Rescaled difference image.

and window dimensions must be selected for the Gaussian smoothing functions. The size of the variance will determine the extent of the deformation and its region of influence. In controlled environments, these variances can be manually selected for good results, as was the case in this paper. However, for situations where the amount of deformation involved and the spacing of control points in the block matching are unknown, then variance selection can have a much greater impact on final results.

From the results it is concluded that the thin-plate spline and hybrid MI-Fluid approach would be appropriate for multimodal applications that require a coarse-to-medium registration. However, these two methods do not rely so heavily on parameter selection. This suggests that these two methods may be the optimal choice in unknown situations. Overall though, if conditions are known, then Gaussian convolution is the better selection as it can be tailored to a situation, is simpler and also computationally faster than both the thin-plate spline and the MI-Fluid approaches.

References

- [1] F. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, June 1989.
- [2] G. Christensen. *Deformable Shape Models for Anatomy*. PhD thesis, Washington University, August 1994.
- [3] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multi-modality image registration based on information theory. In Y. B. et. al., editor, *Information Processing in Medical Imaging*, pages 263–274, The Netherlands, 1995. Kluwer Academic Publishers.
- [4] C. Fookes and M. Bennamoun. Rigid and non-rigid image registration and its association with mutual information: A review. Technical Report ISBN: 1 86435 569 7, RCCVA, QUT, Brisbane, Australia, May 2002.
- [5] T. Gaens, F. Maes, D. Vandermeulen, and P. Suetens. Non-rigid multimodal image registration using mutual information. In *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 1099–1106, 1998.
- [6] T. Gaens, S. Sagaert, and D. Vandermeulen. Non-rigid registration using mutual information. Technical Report 97/10, Katholieke Universiteit Leuven, Medical Image Computing, Leuven, Belgium, 1997.
- [7] N. Hata, T. Dohi, S. Warfield, W. Wells, R. Kikinis, and F. Jolesz. Multimodality deformable registration of pre- and intraoperative images for mri-guided brain surgery. In *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 1067–1074, 1998.
- [8] F. Maes, A. Collignon, D. Vandermeulen, and G. M. P. Suetens. Multimodality image registration by maximisation of mutual information. *IEEE Transactions on Medical Imaging*, 16(2):187–198, April 1997.
- [9] I. Vajda. *Theory of statistical inference and information*. Kluwer, Dordrecht, The Netherlands, 1989.
- [10] P. Viola and W. Wells. Alignment by maximisation of mutual information. *5th International Conference on Computer Vision*, pages 16–23, 1995.




Figure 3. Registration results. Letters (a), (b), and (c) represent results computed with Gaussian convolution, thin-plate spline warp, and MI-Fluid algorithm respectively. Numbers (1), (2), and (3) represent the final image after registration, the rescaled difference image, and a histogram of the intensity differences.