

Intensity-Based Non-rigid Registration Using Adaptive Multilevel Free-Form Deformation with an Incompressibility Constraint

Torsten Rohlfing and Calvin R. Maurer, Jr.

Image Guidance Laboratories, Department of Neurosurgery, Stanford University Medical Center, 300 Pasteur Drive, Room S-012, Stanford, CA 94305-5327, USA
{torsten.rohlfing,calvin.maurer}@igl.stanford.edu

Abstract. A major problem with non-rigid image registration techniques in many applications is their tendency to reduce the volume of contrast-enhancing structures [10]. Contrast enhancement is an intensity inconsistency, which is precisely what intensity-based registration algorithms are designed to minimize. Therefore, contrast-enhanced structures typically shrink substantially during registration, which affects the use of the resulting transformation for volumetric analysis, image subtraction, and multispectral classification. A common approach to address this problem is to constrain the deformation. In this paper we present a novel incompressibility constraint approach that is based on the Jacobian determinant of the deformation and can be computed rapidly. We apply our intensity-based non-rigid registration algorithm with this incompressibility constraint to two clinical applications (MR mammography, CT-DSA) and demonstrate that it produces high-quality deformations (as judged by visual assessment) while preserving the volume of contrast-enhanced structures.

1 Introduction

Non-rigid image registration algorithms [1,5,7] based on free-form deformations [3] have recently been shown to be a valuable tool in various medical image processing applications. However, a major problem with existing algorithms is that when they are applied to pre- and post-contrast image pairs, they produce transformations that substantially shrink the volume of contrast-enhancing structures. Tanner et al. [10] observed this for contrast-enhanced breast lesions and we have observed this for contrast-enhanced vessels in CT-DSA. This problem severely affects the usefulness of the resulting transformation for volumetric analysis, image subtraction, and multispectral classification.

To address this problem, an additional “energy” term is typically added to the intensity-based similarity measure to constrain the deformation to be smooth. Some authors use a mechanically motivated energy term that constrains the “bending energy” of the deformation [7]. A different approach uses “coupled” control points of the deformation to make the contrast-enhanced lesion locally

rigid [10]. This approach of course requires identification of these structures prior to or during registration. Also, it prevents deformation of these structures even in cases where they have actually deformed.

In this paper, we present a novel incompressibility (local volume preservation) constraint based on the Jacobian determinant of the deformation. Soft tissue in the human body is generally incompressible. By penalizing deviations of the local Jacobian determinant of the deformation from unity, this knowledge can be incorporated into the registration process. The Jacobian has been used for analyzing non-rigid transformations [2,10]. This paper, as far as we are aware, is the first report of the use of this mathematical tool to enforce tissue incompressibility during the registration. We also introduce a locally adaptive deformation refinement and present some modifications to the implementation details of existing non-rigid registration algorithms. Both improvements lead to a substantial increase in computational efficiency.

We apply our intensity-based non-rigid registration algorithm using constrained adaptive multilevel free-form deformations to sample MR mammography and CT-DSA images and demonstrate that it produces high quality deformations while preserving the volume of contrast-enhanced structures.

2 Methods

General Registration Algorithm. Our rigid registration algorithm is based on an independent implementation [5,6] of a technique for rigid and affine registration described in Ref. [8]. It uses “normalized mutual information” (NMI) as the similarity measure [9]. In the first step, this method is employed directly for finding an initial rigid transformation to capture the global displacement of both images. The rigid transformation is then used as the initial estimate for the non-rigid registration. The non-rigid algorithm is an independent and modified implementation of a technique presented by Rueckert et al. [7]. It uses the same NMI similarity measure as the rigid registration. However, a different optimization technique (modified steepest-ascent line search [5]) is used to address the problem of the high dimensionality of the search space in the non-rigid case. In addition to the NMI similarity measure E_{NMI} , our technique incorporates an additional penalty term E_{Jacobian} to constrain the deformation of the coordinate space. A user-defined weighting factor ω controls the relative influence of E_{NMI} and E_{Jacobian} , combining both into the overall cost function E_{Total} as follows:

$$E_{\text{Total}} = (1 - \omega)E_{\text{NMI}} + \omega E_{\text{Jacobian}}. \quad (1)$$

A description of the term E_{Jacobian} and its computation are provided below.

B-Spline Deformation. The non-rigid registration algorithm determines the set of parameters of a deformation \mathbf{T} that optimizes the cost function in Eq. 1. \mathbf{T} is defined on a uniformly spaced grid Φ of control points $\phi_{i,j,k}$, where $-1 \leq i < n_x - 1$, $-1 \leq j < n_y - 1$, and $-1 \leq k < n_z - 1$. Control points with i , j , or k equal to either 0 or $n_x - 3$ ($n_y - 3$ and $n_z - 3$ for j and k) are located on the edge of the image data. The spacings between the control points in x -,

y -, and z -directions are denoted by δ_x , δ_y , and δ_z , respectively. For any location (x, y, z) , the deformation at these coordinates is computed from the positions of the surrounding $4 \times 4 \times 4$ neighborhood of control points:

$$\mathbf{T}(x, y, z) = \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 B_l(u)B_m(v)B_n(w)\phi_{i+l,j+m,k+n}. \quad (2)$$

Here, i , j , and k denote the index of the control point cell containing (x, y, z) , and u , v , and w are the relative positions of (x, y, z) in the three dimensions. The functions B_0 through B_3 are the approximating third-order spline polynomials as described by Ref. [3].

The degrees of freedom of a B-spline based transformation are the coordinates of the control points $\phi_{i,j,k}$. In comparison to previous algorithms, ours does not consider these vectors to be offsets from the original control point positions with initial displacements of 0. Instead, we use them as absolute positions and initialize them as

$$\phi_{i,j,k}^{(0)} := (i\delta_x, j\delta_y, k\delta_z). \quad (3)$$

Thus, the actual coordinate transformation is $(x, y, z) \mapsto \mathbf{T}(x, y, z)$ rather than $(x, y, z) \mapsto (x, y, z) + \mathbf{T}(x, y, z)$. This makes application of the actual deformation computationally more efficient for two reasons: First of all, it immediately reduces the number of required real-value additions per transformed coordinate by 3. More importantly, \mathbf{T} is linear with respect to the $\phi_{i,j,k}$. Therefore, an initial rigid transformation A can be incorporated by applying it to the control point positions ϕ :

$$(A \circ \mathbf{T})(x, y, z) = \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 B_l(u)B_m(v)B_n(w)(A\phi_{i+l,j+m,k+n}). \quad (4)$$

As a consequence, the computational cost for every transformed vector is reduced by the cost of a vector-matrix multiplication plus another vector addition. By implicitly applying the rigid transformation *after* the deformation, we ensure that this is mathematically equivalent to explicitly applying A (because computation of i , j , k , u , v , and w remains independent of A).

Deformation Constraint. The design of our deformation constraint is motivated by the observation that most tissues in the human body are incompressible. In a small neighborhood of the point (x, y, z) , the local compression or expansion caused by the deformation \mathbf{T} can be calculated by means of the Jacobian determinant:

$$J_{\mathbf{T}}(x, y, z) = \det \begin{pmatrix} \frac{\partial}{\partial x} \mathbf{T}_x(x, y, z) & \frac{\partial}{\partial y} \mathbf{T}_x(x, y, z) & \frac{\partial}{\partial z} \mathbf{T}_x(x, y, z) \\ \frac{\partial}{\partial x} \mathbf{T}_y(x, y, z) & \frac{\partial}{\partial y} \mathbf{T}_y(x, y, z) & \frac{\partial}{\partial z} \mathbf{T}_y(x, y, z) \\ \frac{\partial}{\partial x} \mathbf{T}_z(x, y, z) & \frac{\partial}{\partial y} \mathbf{T}_z(x, y, z) & \frac{\partial}{\partial z} \mathbf{T}_z(x, y, z) \end{pmatrix}. \quad (5)$$

The value of $J_{\mathbf{T}}(x, y, z)$ is equal to 1 if the deformation at (x, y, z) is incompressible, greater than 1 if there is local expansion, and less than 1 if there

is compression. Since the 3-D spline is the tensor product of independent 1-D functions, its derivative with respect to x can easily be computed as follows:

$$\frac{\partial}{\partial x} \mathbf{T}_x(x, y, z) = \frac{1}{\delta_x} \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 \left(\frac{d}{du} B_l(u) \right) B_m(v) B_n(w) \phi_{i+l, j+m, k+n}. \quad (6)$$

The remaining derivatives are obvious. Computation of the entries of $J_{\mathbf{T}}$ is in fact very similar to computing \mathbf{T} itself. Depending on the position in the matrix, the spline polynomial B in the respective dimension is simply replaced by its derivative. Using the above definition of $J_{\mathbf{T}}$, the incompressibility constraint penalty is defined as the normalized sum of the absolute logarithm of the local scaling divided by the global one:

$$E_{\text{Jacobian}} = \frac{1}{N_{\Phi}} \sum_{i,j,k} \left| \log(J_{\mathbf{T}}(\phi_{i,j,k}^{(0)})/S) \right|, \quad (7)$$

where $S = s_x s_y s_z$ is the total scaling of the initial affine (rigid plus anisotropic scaling) transformation A and s_x , s_y , and s_z are the scale factors in each dimension. This term penalizes deviations of $J_{\mathbf{T}}/S$ from 1, i.e., it penalizes local tissue expansion and compression. The motivation for using the log of the scaling ratio is to equally penalize a local scaling by cS and by S/c and thus achieve equal weighting of local expansion and compression. The penalty term is furthermore normalized by dividing by the number N_{Φ} of control points. This is necessary as during the application of a multilevel strategy the number of control points and therefore the number of addends in 7 increases by a factor of 8 per level.

Two properties of the penalty term E_{Jacobian} should be noted. First, the penalty term does not penalize rigid transformations with anisotropic scaling. Second, as $J_{\mathbf{T}}$ is evaluated only at the control points, the spline coefficients B as well as their derivatives take on particularly simple forms. For this reason, E_{Jacobian} can be computed very efficiently, which is important to keep execution time reasonable. The same is true for the discrete approximation to the derivatives of E_{Jacobian} with respect to the coordinates of any control point. For finite difference approximation to this derivative, a small offset ϵ is added to and subtracted from the respective parameter. E_{Jacobian} is then re-evaluated, yielding E_{Jacobian}^+ and E_{Jacobian}^- . As all addends in 7 outside a $4 \times 4 \times 4$ neighborhood of the affected control point are constant, they cancel out when $\Delta E_{\text{Jacobian}} = \frac{E_{\text{Jacobian}}^+ - E_{\text{Jacobian}}^-}{2\epsilon}$ is calculated. Computation can thus be restricted to the interior of the local neighborhood region.

Multilevel B-Splines. An arbitrary B-spline deformation can be refined to an identical deformation with the control point spacing in every dimension divided by 2. In the 1-D case, the transformed control point positions ϕ'_i of the refined grid can be computed from the coarse control points ϕ_i :

$$\phi'_{2i+1} = \frac{1}{2} (\phi_i + \phi_{i+1}) \quad \text{and} \quad \phi'_{2i} = \frac{1}{8} (\phi_{i-1} + \phi_{i+1} + 6\phi_i). \quad (8)$$

This can be generalized to 3-D by applying the tensor product. For 2-D, the resulting explicit assignments are given in Ref. [3]; for 3-D they can be found in




Fig. 1. Local (*left*) and global (*right*) refinement of the control point grid. Active control points are marked by \bullet in the right picture.

Refs. [4,5]. However, it is sufficiently efficient and considerably more convenient to compute the tensor product on-the-fly using 8 rather than by applying the closed forms.

Grid Refinement and Fixed Control Points. Refinement of the control point grid by a factor of 2 per dimension increases the number of control points by roughly a factor of 8. This leads to a dramatic increase in computation time. However, not everywhere in the deformed image is the refined grid actually helpful. One approach is to refine the grid only in areas where small deformations could not be modeled otherwise (Fig. 1, left). However, for B-spline deformations it is difficult to preserve consistency at the transition from refined to unrefined grid cells.

We have therefore chosen an alternative approach that unites the power of multilevel deformations with the computational efficiency of locally refined control point grids. In fact, the deformation is globally refined. However, control points where there is insufficient local intensity variation (information) to meaningfully estimate finer deformation (e.g., control points in the image background) are fixed to their current positions. This means that their coordinates are no longer considered degrees of freedom of the deformation. Thus, they can be excluded from gradient computation and optimization. The only penalty still arising from the global refinement is the increased amount of memory required to store the locations of all control points in the refined grid. Computational efficiency is not affected by this as computation of the deformation is independent of the number of control points. In fact, as the globally refined grid is more regular than a locally refined one, the globally refined deformation can potentially be computed more efficiently than one defined on a locally refined grid.

Local Entropy. In order to determine which control points are fixed and which are considered variable, a local entropy criterion is applied. For each control point $\phi_{i,j,k}$, the entropy $H_{i,j,k}^R$ of the reference image data is computed over the region $D_{i,j,k}$ that is influenced by its position:

$$H_{i,j,k}^R = \frac{1}{|D_{i,j,k}|} \sum_{\mathbf{x} \in D_{i,j,k}} p(r(\mathbf{x})) \log p(r(\mathbf{x})). \quad (9)$$




Fig. 2. Illustrative example of local entropy criterion. *Left:* Original CT image. *Center:* Local entropy in $4 \times 4 \times 4$ neighborhoods of the B-spline control points. *Right:* Binary image where feature voxels represent neighborhoods where local entropy is greater than 50% of the maximum. Control points where local entropy is less than 50% of the maximum (background voxels in the binary image) are fixed to their current positions.

Here, $r(\mathbf{x})$ are the voxels in $D_{i,j,k}$ and $p(r)$ is the probability of the respective grey value computed over all voxels in $D_{i,j,k}$. In addition, this local entropy is also computed for the corresponding region in the model image under the current deformation. Then, the maximum entropy over all control points is determined separately for the reference and the model image. Using these limits, a control point is fixed if and only if the local entropies associated with it are less than 50% of the respective maximum in the reference and model image. An example of this process is illustrated in Fig. 2. By considering the local entropies from both reference and model, it is ensured that the deformation remains active even in areas where one of the images has little structure as long as there is sufficient information in the other one.

3 Results

We applied the algorithm presented in this paper to two CT-DSA (head-and-neck and abdomen) and one MR mammography image data sets. Both types of study involve acquiring images before and after contrast injection. We registered each pair of images using rigid registration, unconstrained deformation, and deformation with the incompressibility constraint, and then computed subtraction images (post- minus pre-contrast). Figure 3 illustrates the results achieved using our technique in a subtraction image from one of the MR mammography data sets. The deformation algorithm started with a 40 mm control point spacing that was refined to 20 mm and then to 10 mm.

There is considerable artifact in the subtraction images obtained with rigid registration due to tissue deformation not captured by the rigid transformation. The unconstrained and the constrained non-rigid registrations produced subtraction images with substantially reduced motion artifact. Both non-rigid registration subtraction images visually appeared virtually identical, except that




Fig. 3. *From left to right:* Original breast MR image (post contrast), subtraction image (post- minus pre-contrast) after rigid registration, after unconstrained deformation, and after deformation with the incompressibility constraint. The subtraction images are displayed with an inverse gray scale to produce a clear presentation in print. [Image data provided by Michael Jacobs and David Bluemke, The Johns Hopkins University School of Medicine, Baltimore, MD.]

contrast-enhanced structures (vessels in CT-DSA, tumor in MR mammography) visibly shrank in the unconstrained deformation subtraction images, but not in those obtained with constrained deformation. The shrinkage was quantified in the MR mammography study.

Volume Preservation. In the MR mammography image data set, the contrast-enhanced tumor was segmented (by region growing) in the stack of subtraction images used to generate the slices in Fig. 3. This allowed for the volumetric analysis of the different results. We found that after rigid registration the tumor occupied a volume of 4.46 ml. After unconstrained deformation, the tumor volume shrank by 27% to 3.25 ml. After deformation using the incompressibility constraint with a weight of 5×10^{-2} , the tumor volume remained virtually unchanged at 4.45 ml while still achieving the same visual artifact reduction. Using Jacobian weights of 10^{-3} and 10^{-2} , the tumor volume shrank only slightly to 4.31 ml. This reduction of 3% is within the estimated accuracy of the volume computation.

Computational Efficiency. Adaptive fixing of grid control points has two beneficial effects on the computation time. First, fixed parameters need not be considered for gradient estimation, thus directly reducing the computation time per optimization step. Also, the effective dimension of the search space is reduced, which decreases the number of optimization steps required to find the final solution. For typical image data, we experienced the fraction of fixed parameters to be somewhere between 30% and 60% depending on the ratio of object and background voxels. This led to an observed reduction of computation time by at least 50%.

4 Discussion

We have presented a non-rigid registration algorithm with a novel incompressibility constraint. The Jacobian-based penalty term is effective and efficient. It

is relatively insensitive to the choice of the weighting factor with respect to the intensity-based similarity measure. This is essential since the two terms in our cost function (NMI and Jacobian-based penalty term) are fundamentally different entities and thus there is no a-priori “correct” weighting.

Preliminary results on three image data sets suggest that incorporation of this constraint improves non-rigid registration of pre- and post-contrast images by substantially reducing the problem of shrinkage of contrast-enhanced structures. The energy term as well as its discrete derivative is easy to implement and can be computed rapidly. In fact, almost no performance penalty is experienced when using the Jacobian-based constraint. Our algorithm also incorporates locally adaptive multilevel refinement, which improves computational efficiency of the non-rigid registration.

5 Acknowledgments

The authors began this research in the Department of Neurological Surgery, University of Rochester, Rochester, NY. They gratefully acknowledge support for this work provided by the Ronald L. Bittner Endowed Fund in Biomedical Research. The authors thank Michael Jacobs and David Bluemke of The Johns Hopkins University School of Medicine, Baltimore, MD, for providing the image data shown in Figure 3.

References

1. ERE Denton, LI Sonoda, D Rueckert, DLG Hill, et al. Comparison and evaluation of rigid, affine, and nonrigid registration of breast MR images. *J Comput Assist Tomogr*, 23(5):800–805, 1999.
2. T Hartkens, D Hill, C Maurer, Jr, A Martin, et al. Quantifying the intraoperative brain deformation using interventional MR imaging. *Proc Int Soc Magn Reson Medicine*, 8:51, 2000.
3. S Lee, G Wolberg, SY Shin. Scattered data interpolation with multilevel B-splines. *IEEE Trans Visualization Comput Graphics*, 3(3):228–244, 1997.
4. T Rohlfing. *Multimodale Datenfusion für die bildgesteuerte Neurochirurgie und Strahlentherapie*. PhD thesis, Technische Universität Berlin, 2000.
5. T Rohlfing, CR Maurer, Jr, WG O’Dell, J Zhong. Modeling liver motion and deformation during the respiratory cycle using intensity-based free-form registration of gated MR images. In *Medical Imaging: Visualization, Display, and Image-Guided Procedures*, volume 4319. Proceedings of SPIE, 2001. (In press.)
6. T Rohlfing, JB West, J Beier, T Liebig, et al. Registration of functional and anatomical MRI: Accuracy assessment and application in navigated neurosurgery. *Comput Aided Surg*, 5(6):414–425, 2000.
7. D Rueckert, LI Sonoda, C Hayes, DLG Hill, et al. Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Trans Med Imaging*, 18(8):712–721, 1999.
8. C Studholme, DLG Hill, DJ Hawkes. Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Med Phys*, 24(1):25–35, 1997.

9. C Studholme, DLG Hill, DJ Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognit*, 33(1):71–86, 1999.
10. C Tanner, JA Schnabel, D Chung, MJ Clarkson, et al. Volume and shape preservation of enhancing lesions when applying non-rigid registration to a time series of contrast enhancing MR breast images. In *MICCAI 2000*, pages 327–337.