

Methods for Inverting Dense Displacement Fields: Evaluation in Brain Image Registration*

William R. Crum¹, Oscar Camara¹, and David J. Hawkes¹

Centre for Medical Image Computing, University College London, UK
b.crum@ucl.ac.uk

Abstract. In medical image analysis there is frequently a need to invert dense displacement fields which map one image space to another. In this paper we describe inversion techniques and determine their accuracy in the context of 18 inter-subject brain image registrations. Scattered data interpolation (SDI) is used to initialise locally and globally consistent iterative techniques. The inverse-consistency error, E_{IC} is computed over the whole image and over 10 specific brain regions. SDI produced good results with mean (max) $E_{IC} \sim 0.02\text{mm}$ (2.0mm). Both iterative method produced mean errors of $\sim 0.005\text{mm}$ but the globally consistent method resulted in a smaller maximum error (1.9mm compared with 1.4mm). The largest errors were in the cerebral cortex with large outlier errors in the ventricles. Simple iterative techniques are, on this evidence, able to produce reasonable estimates of inverse displacement fields provided there is good initialisation.

1 Introduction

In most pair-wise non-rigid registration applications, either by design or because of the available registration software, there is an implied direction i.e. one image is designated the target (reference) and the other is the source (floating image). The transformation, \mathbf{T} , is defined at points \mathbf{t} in the target space and specifies the point in source space $\mathbf{s} = \mathbf{t} + \mathbf{T}(\mathbf{t})$ corresponding to each target voxel. The transformation may be parameterised at a coarser level but voxel-wise displacements can always be obtained by interpolation. Often the choice of target and source is application-specific e.g. to propagate a mesh from one image to another requires the mesh to be defined in the target space but to propagate a set of dense labels requires the labels to be defined in the source space. Similarly, to create a group average template requires the template space to be the target, but to subsequently transform labels from the template space to individual scans requires the template to be the source. The inverse transformation (from source to target) cannot be trivially deduced from the forward transformation. A registration algorithm can be run in reverse (i.e. with target

* This work is supported by the EPSRC Integrated Brain Image Modelling, Medical Image and Signals IRC and Modelling and Understanding and Predicting Structural Brain Change projects We are thankful to David Kennedy, Centre for Morphometric Analysis, MGH, Boston for the brain images used in this work.

and source exchanged) but the reverse transformation is not generally consistent with the forward transformation.

Some registration algorithms are designed to be inverse-consistent [1][2] and compute consistent forward and backward transformations. The related symmetric approaches [3] are conceptually similar but improve robustness without computing explicit inverses. In theory, inverse-consistent algorithms should remove the need to worry about the reverse transformation. In practice, these algorithms are not all that mature or widely available and the effect of enforcing strict inverse-consistency on the registration optimisation can be computationally expensive and has not been studied in any detail. In addition, it is clear that for many applications strict inverse-consistency is not appropriate. For instance, registering scans of a single patient anatomy subject to some mechanical deformation is a case where inverse-consistency is appropriate. However registering inter-subject brain scans is a case where strict inverse consistency has little biological foundation [4].

Some study of the computation of inverse transformations has been made in the registration literature, most notably in [1] and [5]. In [6] methods for making a discrete transformation topology-preserving, and therefore invertible, are described which could be embedded in a registration algorithm. However, to date there has not been a systematic evaluation of numerical methods for inversion of non-rigid registration transformations. Therefore, in this paper we evaluate different methods for computing the inverse of a dense, non-parametric, displacement field. We make no assumptions beyond assuming that the Jacobian determinant remains positive and comment on exceptions later. We focus on two classes of techniques (i) scattered data interpolation (SDI) and (ii) inverse-consistent iterative approaches. We evaluate the accuracy of these techniques globally and over specific neuroanatomical regions for transformations obtained from inter-subject MR brain registration.

2 Method

2.1 Inversion Problem

The inversion problem is sketched in figure 1. We assume the forward transformation is known at the target voxel centres, \mathbf{t} , where solid arrows point to corresponding points in the source space. To invert this transformation the magnitude and direction of the dashed arrows must be deduced at voxel-centres, \mathbf{s} , in the source space.

Inversion Using Direct Interpolation. The simplest interpolated inversion is a nearest-neighbour (NN) approach where the inverse transformation is defined as the negative nearest forward transformation vector. We adopted an extended NN method (NNe) as the simplest initialiser for the iterative methods described below. Where the nearest forward transformed point did not lie within the current source voxel because of local divergence, an average of forward transformed points in surrounding voxels was used. This was computed hierarchically so that




Fig. 1. Left: The forward transformation is defined in the target space at \mathbf{t} . Right: The inverse transformation must be deduced in the source space at \mathbf{s} . Filled dots indicate where the inverse transformation is defined by the forward transformation. The dashed circle indicates a possible search space for a scattered data interpolation approach.

averages with the highest number of nearest-neighbours were computed before averages with a smaller number of neighbours.

Inversion can also be cast as a simple scattered data interpolation (SDI) problem. In figure 1, the inverse transformation is known by definition at the points in the source image pointed to by the forward transformation and can be deduced at the voxel centres \mathbf{s} by interpolation [7]. A distance weighting function, $w()$ sets the relative contribution of each scattered point to the interpolated inverse displacement; we chose an inverse-square (Shepard) weighting function (equation 1) where d_i is the distance from the interpolated point to the i^{th} data point.

$$\mathbf{I}(\mathbf{s}) = \frac{\sum_i w(d_i) * \mathbf{I}(\mathbf{t}_i + \mathbf{T}_i)}{\sum_i w(d_i)} \tag{1}$$

As only local points contribute significantly, we use a modified weighting function incorporating a search radius, R , given by $w(d_i; R) = \left(\frac{1}{d_i} - \frac{1}{R}\right)^2$ when $d_i \leq R$ and $w(d_i; R) = 0$ otherwise. R is initially chosen to be twice the maximum voxel dimension but if no points are found it is arbitrarily increased until at least 4 scattered points are included.

Inversion Using Iterative Techniques. We compared two iterative techniques, one inverting point-by-point iteratively in target space (the **local** method) and one inverting the entire field iteratively in source space (the **global** method).

The **local** method is implemented as described in [1] and searches for the position in target-space, \mathbf{t} where the trilinearly interpolated forward transformation,

```

FOREACH VOXEL,  $\mathbf{t}$ 
   $i = 0$ 
  DO
     $\mathbf{R}_L = \mathbf{t} + \mathbf{T}(\mathbf{t}) - \mathbf{s}$ 
     $\mathbf{t}^{i+1} = \mathbf{t}^i - \mathbf{R}_L/2$ 
     $i = i + 1$ 
  UNTIL ( $\mathbf{R}_L \leq \text{tol}$ ) OR ( $i > 200$ )
   $\mathbf{I}(\mathbf{t}) = -\mathbf{T}(\mathbf{t}^{i-1})$ 
NEXT VOXEL

```

Fig. 2. Pseudo-code for **local** iterative inversion

$\mathbf{t} + \mathbf{T}(\mathbf{t})$ points to the current source voxel centre, \mathbf{s} . A residual error \mathbf{R}_L is used to update the current estimate of \mathbf{t} at each voxel. Figure 2 shows pseudo-code for the **local** method.

The **global** method uses the current estimate of the inverse transformation \mathbf{I} to interpolate (trilinearly) the forward transformation \mathbf{T} to each source voxel-centre, \mathbf{s} . Then the residual difference between the estimated inverse and the interpolated forward transformation is used to update the inverse as before. The magnitude of the update is capped for stability and each iteration is computed over the whole field before the next iteration. Figure 3 shows pseudo-code for the **global** method.

Both methods have a maximum number of iterations (200) specified and a stopping criterion based on residual size. The stopping criterion (residual $< 0.01\text{mm}$) for the **local** method is inherently stricter than for the **global** method as it is computed on a point-by-point basis rather than averaged over the volume. For this reason the **global** method has an additional stopping criterion based on the maximum inverse-consistency error being $\leq 0.1\text{mm}$.

```

 $i = 0$ 
DO
   $max = av = 0$ 
  INTERPOLATE  $\mathbf{T}(\mathbf{s})$  FROM  $\mathbf{T}(\mathbf{t})$  USING  $\mathbf{I}()$ 
  FOREACH VOXEL,  $\mathbf{t}$ 
     $\mathbf{R}_G(\mathbf{s}) = \mathbf{I}(\mathbf{s}) + \mathbf{T}(\mathbf{s})$ 
     $\mathbf{I}^{i+1}(\mathbf{s}) = \mathbf{I}^i(\mathbf{s}) - \alpha \mathbf{R}_G(\mathbf{s})$ 
     $av = av + |\mathbf{R}_G|$ 
     $max = \text{MAX}(|\mathbf{R}_G|, max)$ 
     $i = i + 1$ 
  NEXT VOXEL
   $av = av/nvox$ 
UNTIL ( $(av \leq \text{tol})$  AND ( $max \leq \text{tolmax}$ )) OR ( $i > 200$ )

```

Fig. 3. Pseudo-code for **global** iterative inversion

2.2 Image Registration Experiment

Image-data and structural labels from 18 subjects were supplied by the Centre for Morphometric Analysis, Boston. The scans were MRI T1-weighted coronal volume acquisitions of dimensions 256x256x128 and voxel sizes 1.0x1.0x1.5mm. All subjects were of normal appearance (i.e. had no obvious pathology or resection). A standard (i.e. not inverse-consistent) non-rigid (fluid) registration algorithm [8] was used with two multi-resolution steps with a maximum of 200 iterations per step, driven by intensity cross-correlation. The 18 images were grouped randomly into 9 pairs and registered in both directions giving 18 registrations, i.e 9 forward and 9 reverse, so that the inherent inverse-consistency error of the registration algorithm could also be computed. Then the inverse of each registration transformation was computed using (a) SDI (b) NNe+local (c) SDI+local (d) NNe+global (e) SDI+global.

$$E_{IC} = |\mathbf{t}' - \mathbf{t}| \text{ where } \mathbf{t}' = \mathbf{t} + T(\mathbf{t}) + I(\mathbf{t} + T(\mathbf{t})) \tag{2}$$

We evaluated the different inversion methods by computing the inverse-consistency error (equation 2) at each point and compared this with the inherent registration inverse-consistency in corresponding forward and reverse registrations. We also computed the inversion error over ten specific brain structures (amygdala, caudate, cerebellum, cortex, hippocampus, lateral ventricles, pallidum, putamen, thalamus, white-matter). To assess how the error scaled with transformation magnitude a single transformation from the test set was (a) composed with global rotational components between 2.5° and 25.0° (b) globally scaled by factors in the range 1.25 to 5.00, and inverted as above.

3 Results

Figure 4 shows the mean and maximum global inverse-consistency error for the five inversion techniques. SDI produces sub-voxel inverse-consistency on average




Fig. 4. The mean (left) and maximum (right) inverse-consistency error over 18 inter-subject brain registrations. S=Scattered Data Interpolation, N = Nearest-Neighbour-extended, L = Local Iterative and G = Global Iterative.

and using SDI to initialise either iterative method reduces the mean error further by $\sim 75\%$. In this experiment the SDI+**local** and SDI+**global** iterative methods are virtually indistinguishable as regards mean error but there is some regional variation apparent from maps of E_{IC} (figure 5). The SDI+**global** method had a lower maximum error averaged over all the registrations but the SDI+**local** method had a lower worst-case maximum error. By comparison, the inverse-consistency error determined directly from the forward and reverse non-rigid registration was much larger (0.7mm compared with 0.02mm for the SDI case).




Fig. 5. Example of an inverse consistency error map for a single subject. Left Panel = subject anatomy. The left-hand-side of the middle and right panels shows the error distribution for S+L and S+G inversion methods respectively ranging from 0.0mm=black ≥ 0.6 mm=white. For comparison the error distribution for fwd-rev registration is shown on the right-hand-side scaled from 0.0mm = black ≥ 6.0 mm = white.

Figure 6 shows that the mean inversion error for all brain structures was $\ll 0.1$ mm and that with the exception of the cerebral cortex (where the **global** error had significantly smaller mean error), the performance of **local** and **global** techniques is virtually identical. These errors should be considered in the context of the mean (maximum) displacements computed over all registrations for all structures of 4.2 (20.7)mm. We expect most inversion problems in the cortex region as inter-subject variation results in particularly tortuous non-rigid transformations. Figure 6 also shows the maximum inversion error over all registrations associated with each label. Overall the **global** inverses have lower maxima than the **local** with the most notable exception being in the hippocampus. Most structures have maximum errors ≤ 1 mm, but the cortex, lateral ventricles and white matter structures have maximum errors ~ 3 mm. The ventricle results are dominated by two outlier cases with mean errors 10 times the rest of the cohort. Excluding these outliers gives a mean (maximum) error of 0.03 (0.08)mm for ventricles for both **local** and **global** methods. More detailed analysis is required to understand these cases which do not exhibit significantly larger deformation magnitudes than the rest of the cohort or have ventricle errors correlated with errors in the cortex. Figure 7 shows how the inversion error increases with the maximum rotational or scaled displacement. Note that the rotational displacement, which preserves local volume change, has a shallow linear error relationship




Fig. 6. The mean (left) and maximum (right) inverse-consistency error over 18 inter-subject brain registrations. am=amygdala, ca=caudate, ce=cerebellum, co=cortex, hp=hippocampus, lv = lateral ventricles, pa=pallidum, pu=putamen, th=thalamus, wm=white-matter.

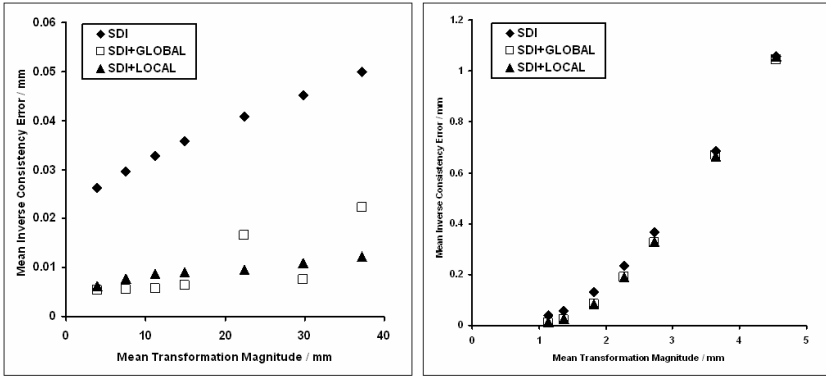


Fig. 7. The scaling of the mean inverse-consistency error with applied (left) rotation and (right) scaling

but the scaled displacement, which results in negative Jacobian determinants, has a polynomial error relationship.

4 Discussion

We have evaluated simple techniques for inversion of dense displacement fields in the context of inter-subject non-rigid registration algorithms and quantified the inversion error over a range of neuroanatomical structures. Both locally and globally consistent techniques initialised with a scattered data interpolation approach proved generally reliable and accurate in this study. We confirmed that inverting transformations matching convoluted structures such as the cortex is prone to relatively higher errors. While this is not unexpected it suggests

that inverse-consistent registration approaches which rely on an explicit inverse calculation to enforce inverse-consistency will also perform less well in these regions.

All the original registration transformations were diffeomorphic i.e. the Jacobian determinant $|J| > 0$ everywhere and the transformation inverse was therefore well-defined throughout the domain. Many non-rigid algorithms do not guarantee a positive Jacobian determinant and therefore some transformation regularisation or transformation model-based constraints may be necessary.

The methods described here are completely generic and make no assumptions about the classes of permitted displacements. However these methods should be considered the "lowest common denominator" of inversion techniques and more sophisticated approaches will improve robustness. In future work we will analyze regions of displacement fields with high E_{IC} and investigate the use of regularising techniques in the spirit of [6] for regions with $|J| < 0$. These findings will be used to improve the performance of inverse-consistent registration algorithms.

References

1. Christensen, G.E, Johnson, H.J: Consistent image registration. *IEEE Transactions on Medical Imaging* 20(7), 568–582 (2001)
2. Cachier, P., Rey, D.: Symmetrization of the non-rigid registration problem using inversion-invariant energies: application to multiple sclerosis. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 472–481 (2000)
3. Rogelj, P., Kovačič, S.R.: Symmetric image registration. *Medical Image Analysis* 10(3), 484–493 (2006)
4. Crum, W.R., Griffin, L.D., Hill, D.L., Hawkes, D.J: Zen and the art of medical image registration: Correspondence, homology and quality. *NeuroImage* 20, 1425–1437 (2003)
5. Rao, A., Chandrashekar, R., Sanchez-Ortiz, G.I., Mohiaddin, R., Aljabar, P., Hajnal, J.V., Puri, B.K., Rueckert, D.: Spatial transformation of motion and deformation fields using nonrigid registration. *IEEE Transactions on Medical Imaging* 23(9), 1065–1076 (2004)
6. Karacali, B., Davatzikos, C.: Estimating topology preserving and smooth displacement fields. *Journal of Electronic Imaging* 23(7), 868–880 (2004)
7. Amidror, I.: Scattered data interpolation methods for electronic imaging systems: a survey. *Journal of Electronic Imaging* 11(2), 157–176 (2002)
8. Crum, W.R, Tanner, C., Hawkes, D.J: Multiresolution, anisotropic fluid registration: Evaluation in magnetic resonance breast imaging. *Physics in Medicine and Biology* 50, 5153–5174 (2005)