

A Generic Framework for Non-rigid Registration Based on Non-uniform Multi-level Free-Form Deformations

Julia A. Schnabel¹, Daniel Rueckert², Marcel Quist³, Jane M. Blackall¹,
Andy D. Castellano-Smith¹, Thomas Hartkens¹, Graeme P. Penney¹, Walter A. Hall⁴,
Haiying Liu⁵, Charles L. Truwit⁵, Frans A. Gerritsen³, Derek L. G. Hill¹, and
David J. Hawkes¹

¹ Computational Imaging Science Group, Radiological Sciences and Medical Engineering,
Guy's Hospital, King's College London, UK
julia.schnabel@kcl.ac.uk

² Visual Information Processing, Dept. Computing, Imperial College of Science, Technology
and Medicine, London, UK

³ EasyVision Advanced Development, Philips Medical Systems, Best, NL

⁴ Dept. Neurosurgery, University of Minnesota, Minneapolis, MN, USA

⁵ Dept. Radiology, University of Minnesota, Minneapolis, MN, USA

Abstract. This work presents a framework for non-rigid registration which extends and generalizes a previously developed technique by Rueckert et al. [1]. We combine multi-resolution optimization with free-form deformations (FFDs) based on multi-level B-splines to simulate a non-uniform control point distribution. We have applied this to a number of different medical registration tasks to demonstrate its wide applicability, including interventional MRI brain tissue deformation compensation, breathing motion compensation in liver MRI, intra-modality inter-modality registration of pre-operative brain MRI to CT electrode implant data, and inter-subject registration of brain MRI. Our results demonstrate that the new algorithm can successfully register images with an improved performance, while achieving a significant reduction in run-time.

1 Introduction

Non-rigid image registration is playing an increasingly important role in both clinical and research applications. Even though there is a large number of registration methods to be found in the literature [2,3], existing methods have usually been designed for or fine-tuned to specific, mostly single-modality applications. In this paper we propose a flexible framework for non-rigid registration which covers a range of deformation tasks in medical applications. This framework extends and generalizes a previously published method based on free-form deformations (FFDs) using B-splines which was developed for registration of contrast enhanced MR mammography images [1]. The original method was formulated as a two-stage process: first, the global motion is corrected using a rigid or affine transformation. The global motion then becomes the starting estimate for the second stage, where the local motion is further modelled using FFDs based on B-splines. Manipulating the underlying mesh of control points yields a

smooth deformation of structures embedded in the image, where the control points act as parameters of the transformation. The combined motion model can be written as:

$$\mathbf{T}(x, y, z) = \mathbf{T}_{global}(x, y, z) + \mathbf{T}_{local}(x, y, z) \quad (1)$$

with the local motion at each point given by the 3D tensor product of the familiar 1D cubic B-splines [4]. The optimal transformation \mathbf{T} is determined by minimizing a registration cost function:

$$\mathcal{C} = -\mathcal{C}_{similarity}(\mathcal{I}_A, \mathbf{T}(\mathcal{I}_B)) + \lambda \mathcal{C}_{deformation}(\mathbf{T}) \quad (2)$$

The similarity term maximizes the voxel similarity between the image pair, and is chosen to be normalized mutual information (NMI) [5]. The deformation cost term is defined as the 3D equivalent of a thin-plate bending energy in order to maximize the smoothness of the transformation, weighted by a factor λ .

The performance of this registration method is limited by the resolution of the control point mesh, which is linearly related to the computational complexity: more global and intrinsically smooth deformations can only be modelled using a coarse control point spacing, whereas more localized and intrinsically less smooth deformations require a finer spacing. If it is known *a priori* which magnitude of deformation is to be expected, the mesh resolution can be chosen accordingly, and folding can be penalized using a suitable deformation cost term. However, often a range of deformations needs to be modelled which cannot be captured by a single mesh resolution. More importantly, it may be desirable to have a non-uniform control point spacing to restrict deformation to localized regions in the image pair, while excluding regions where the images are already in alignment, form part of image background, or have been identified as rigid bodies. For many applications, non-uniform control point spacing may not be necessary, however the computational complexity can be decreased without compromising the registration performance. In the following, we present our multi-resolution non-rigid registration framework that can have a non-uniform control point distribution.

2 Framework

2.1 Multi-resolution Registration

To model a large range of deformations, a multi-resolution mesh representation is needed. This issue has been addressed in [1] in a coarse-to-fine fashion where the mesh is progressively refined by alternating between mesh deformation and mesh subdivision using suitable B-spline subdivision techniques [6]. As an alternative, Rueckert et al. [1] formulated an approach using multi-level B-splines to create a hierarchy of local deformation meshes [4] which will be presented in this work. Multi-resolution registration using this concept is then achieved by deforming a sequence of control point meshes Φ^1, \dots, Φ^H of arbitrarily increasing resolution using multi-level B-splines (each level corresponding to a single-resolution B-spline mesh). After registration, the local deformation of each point in the image volume domain is given by the sum of the local deformations across levels:

$$\mathbf{T}_{local}(x, y, z) = \sum_{h=1}^H \mathbf{T}_{local}^h(x, y, z) \quad (3)$$

where each $\mathbf{T}_{local}^h(x, y, z)$ is computed with respect to the B-spline of that level h . To avoid the overhead of recalculating previously recovered local deformations up to level $h - 1$, these can be efficiently pre-computed for the deformation of level h .

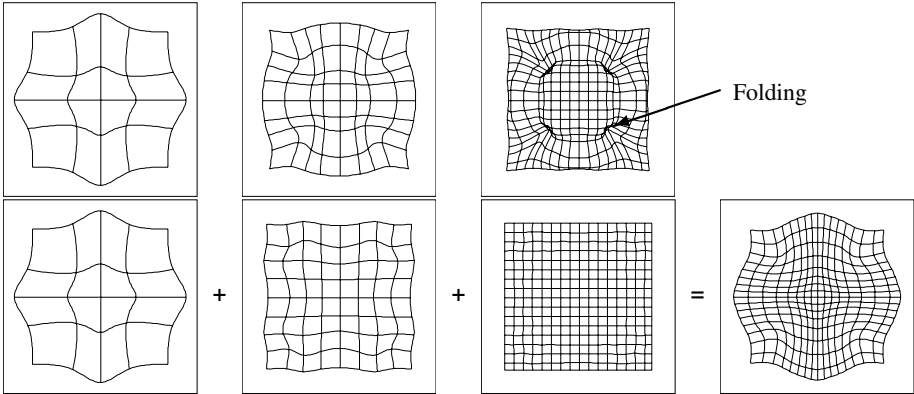


Fig. 1. 2D slices through 3D single- and multi-resolution FFDs for a synthetic registration task. Top: Single-resolution FFDs are limited by low mesh resolutions, and may develop folding at high resolutions (see arrow). Bottom: Multi-resolution FFDs are modelling deformations of different magnitudes at the respective mesh resolution, without developing any folding. Note that at intermediate resolution, deformations of the preceding lower resolution are partially corrected.

Fig. 1 shows 2D slices through 3D single- and multi-resolution FFDs for the registration of a synthetic cube and sphere. One can observe that multi-resolution FFDs can model large deformations without folding at higher resolution levels, whereas FFDs at a single mesh resolution are locally not flexible enough at low resolutions, and can be subject to folding at higher resolutions without any additional smoothness constraints. Hence, when using multi-resolution FFDs, the deformation cost term regularizing the smoothness of the deformation in Eq. 2 is no longer crucial and we can set $\lambda = 0$.

2.2 Non-uniform Control Point Spacing

Non-uniform control point spacing for FFDs is a well known concept from computer graphics. For example, Non-Uniform Rational B-Splines (NURBS) [7] have been used for FFDs to represent and interactively manipulate object shapes. We propose to simulate this concept by extending the multi-resolution registration presented in section 2.1, while avoiding the more complex lattice traversal and control point manipulation of NURBS. We introduce a control point status \mathcal{S} associated with each control point at each level in the multi-resolution mesh hierarchy, marking it as either *active* or *passive*.




Fig. 2. Simulation of non-uniform control point spacing by combining multi-resolution registration with a control point status. Active control points are marked as black, passive ones as grey.

Active control points are allowed to move during the registration process, whereas passive control points remain fixed. Fig. 2 illustrates this concept. Prior knowledge derived from segmentation may be used to assign a control point status. Alternatively, we have developed two different, complementary approaches which derive the status automatically from the image data. Both investigate the local B-spline support regions $\Phi_{i,j,k}$ of each control point $\phi_{i,j,k}$, which provide natural bounds for the local support region that a B-spline control point at a certain resolution level has. Within these bounds, and over the whole image domain, a statistical measure \mathcal{M} can be calculated from the underlying image data to determine the status of each control point as:

$$\mathcal{S}(\phi_{i,j,k}) = \begin{cases} \text{active} & \text{if } \mathcal{M}(\Phi_{i,j,k}) > \epsilon \\ \text{passive} & \text{otherwise} \end{cases} \quad (4)$$

where ϵ is a selection threshold, which may be normalized by the global measure \mathcal{M} . In the following, we distinguish between measures computed over the reference image, and joint measures computed over the image pair.

Reference image measures characterize the reference image locally, and are applicable for serial registration to a baseline scan, inter-subject registration to a common reference subject, or registration to an atlas. They are computed prior to the registration process for all undeformed meshes in the B-spline hierarchy. One such measure which describes the information content of the image intensity distribution is the Shannon-Wiener entropy $H = -\sum p(a) \log p(a)$, where $p(a)$ is the probability that an image voxel has intensity a . Other local intensity measures include intensity variance σ , features based on higher order image differentials, or more complex noise estimators. Using such measures, image background regions can be naturally excluded without explicit segmentation, and local image structure is deformed with the appropriate mesh resolution. For the inter-subject registration application in this work, we have investigated both H and σ as suitable reference image measures.

Joint image pair measures describe the degree of local image alignment. They need to be recomputed for each level after the deformation of the preceding levels for progressive adaptation, and include local measures such as intensity or gradient differences, local correlation, and information-theoretic measures computed over local histograms. These measures should be normalized by the integration measure over the whole image domain. We propose a more consistent, generalized joint image pair measure which is based on the local gradient $\partial\mathcal{C}/\partial\phi_{i,j,k}$ of the global registration cost term of Eq. (2), which does not rely on potentially insufficient statistics defined over small local image

regions or histograms. Computing the gradient of the cost functions to determine local adaptivity has also very recently been proposed in [8].

3 Applications

In the following, we present example results for the new framework for intra- and inter-modality, intra-subject and inter-subject registration tasks. Figs. 3-6 show example 2D slices through 3D volumes, where registration has been performed on 3D volume pairs.

Intra-subject intra-modality registration:




Fig. 3. Intra-subject registration of pre- to post-operative brain MRI. Subtraction images after registration: (a) rigid, and non-uniform non-rigid using joint image pair adaptation at (b) 15mm spacing, and (c) 15mm refined to 7.5mm spacing.

Fig. 3 shows an example slice through subtraction images after rigid and non-rigid registration of an MP-RAGE MR brain scan before dura puncture to an MP-RAGE MR brain scan after a functional surgical procedure where a unilateral thalamic stimulator was inserted stereotactically to suppress tremor. After rigid registration, a shift in the brain tissue can be observed which may have been caused by patient positioning on the operating table, loss of cerebral spinal fluid, and tissue deformation. It was previously shown that the non-rigid algorithm by Rueckert et al. [1] can correct for this shift [9]. The locality of the shift, and the expected magnitude of up to 15mm for this case can be used to adapt B-spline meshes of 15mm and refinement to 7.5mm control point spacing using the local similarity gradient as a joint image pair measure M ($\epsilon = 0.01$). Control points in areas where shift can be observed, and where $M > \epsilon$, are automatically marked as *active*, and as *passive* in remaining areas which are sufficiently aligned after rigid registration. This has reduced the number of active control points and the associated computing time by 60% (15mm) and 97% (7.5mm) for this case.

Fig. 4 shows an example for registration of liver MRI between exhale and inhale positions for a volunteer. To correct for the rigid, mostly translational movement of the




Fig. 4. Intra-subject registration of liver MRI. Inhale contour overlays onto exhale image after (a) rigid and (b) non-rigid registration. Deformation was restricted to the liver at exhale.

liver, and the additional deformation within the liver tissue, we have segmented the liver in the reference image (at exhale) prior to registration to exclude all external structure from the registration process. Non-rigid registration was performed using a mesh resolution of 23.12mm, chosen as a multiple of the in-plane voxel size. Note the improved alignment of the surface and vessels within the liver after non-rigid registration, shown as contour overlays. The non-uniform control point spacing has reduced the computing time by around 94%.

Intra-subject inter-modality registration:




Fig. 5. Intra-subject registration of pre-operative brain MRI to CT with electrode implants. (a) CT, and CT contour overlays showing brain surface and electrodes onto MR after (b) rigid and (c) non-rigid registration. Deformation was restricted to the intracranial cavity.

Fig. 5 shows an example for inter-modality registration for a pre-operative MR brain scan to a post-operative CT scan of a patient with epilepsy after an electrode grid implantation to map electrical activity. Similar to the liver registration task, we have performed a segmentation prior to registration. The reason for this is that the high intensity values of the electrode grid would otherwise be mapped to the dura in the MR scan. Adapting the mesh with respect to setting only the control points within the intra-cranial cavity to *active* reduced the non-rigid registration time by 85%, and, more importantly, the shift of the brain surface caused by the grid implantation was recovered without deforming the dura. Another example for intra-subject inter-modality registration, which is not shown in this paper, is the registration of pre-contrast to contrast-enhanced MR mammographic images. In [1], B-spline subdivision was proposed to perform a coarse-to-fine deformation compensation for that application, which alternatively can be achieved using the multi-resolution scheme presented in this work.

Inter-subject intra-modality registration:




Fig. 6. Inter-subject registration of brain MRI of 7 healthy subjects. ROI of (a) reference subject and averages after registration: (b) affine, (c) uniform non-rigid at 20mm, (d) uniform non-rigid at 20mm, refined to 10mm and 5mm, and non-uniform non-rigid at 20mm, refined to 10mm and 5mm using (e) entropy H and (f) variance σ as reference image measures.

As an example for inter-subject, intra-modality registration, we have registered MR T1-weighted brain volumes of 6 control subjects to one reference control subject. After

registration, averages of the aligned image volumes were generated. Fig. 6 shows an example region of the reference subject and averages after affine and non-rigid registration using uniform control point spacing of 20mm, and refinement via 10mm to 5mm spacing. One can visually perceive the improved alignment when going from affine to non-rigid registration using 20mm spacing, as well as an improved refinement to 5mm. We have also used non-uniform control point spacing on the basis of the local entropy ($\epsilon = 0.75H_{total}$) and variance ($\epsilon = 0.5\sigma_{total}$) as reference image measures. The result is comparable to the uniform mesh refinement, but needed considerably less computing time (reduction by 43-66% for $M = H$ and by 29-55% for $M = \sigma$, for increasing mesh resolutions, respectively). Using a 5mm control point spacing from the start was found to be computationally prohibitively expensive. Both reference image measures were found to be robust and well-behaved estimates with respect to the control point support regions and a measure-specific, but constant threshold ϵ across resolutions.

4 Discussion and Conclusions

We have presented a non-rigid image registration framework based on multi-resolution refinement of adaptable FFDs using B-splines, which extends the work of Rueckert et al. [1]. We have applied the framework to a variety of medical registration problems, demonstrating its flexibility and wide applicability. In particular, its ability to constrain deformation to selective regions, on the basis of segmentation, or automatically using reference or joint image pair measures provides a large gain in computing time without any apparent loss of registration quality. Although folding of the deformation field was not explicitly constrained to ensure diffeomorphism, we have found no occurrences of folding because of the intrinsically smooth deformation modelling capabilities of the multi-resolution FFD. Validation of non-rigid registration is a challenging task, and often limited to visual assessment. We have recently developed a biomechanical deformation simulator using finite element methods [10] which we have successfully applied to the registration algorithm by Rueckert et al. for contrast-enhanced MR mammography [1]. We are planning to use this method to validate and further improve the non-rigid registration framework presented in this paper. In particular, we will investigate the parameter selection for the adaptivity criteria, and the range and sampling values for the multi-level FFDs. For inter-subject registration, we are investigating target registration errors on the basis of anatomical landmarks which is subject to ongoing work.

Acknowledgements

JAS is funded by Philips Medical Systems EV-AD, DR is partially funded by EPSRC GR / N / 24919, GPP is funded by EPSRC GR / M53752, ADCS and TH are funded by EPSRC GR / M47294. The electrode data were provided by Prof. Charles Polkey from Dept. Healthcare, KCL, the brain control data by the Institute of Psychiatry, KCL, the interventional MRI data by the University of Minnesota, and the liver data were acquired at Guy's Hospital.

References

1. D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes. Non-rigid registration using Free-Form Deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.
2. H. Lester and S. R. Arridge. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32(1):129–149, 1999.
3. J. V. Hajnal, D. L. G. Hill, and D. J. Hawkes. *Medical image registration*. CRC Press, 2001.
4. S. Lee, G. Wolberg, and S. Y. Shin. Scattered data interpolation with multilevel B-splines. *IEEE Transactions on Visualization and Computer Graphics*, 3(3):228–244, 1997.
5. C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap entropy measure of 3D medical image alignment. *Pattern Recognition*, 32:71–86, 1999.
6. D. R. Forsley and R. H. Bartels. Hierarchical B-spline refinement. *ACM Transactions on Computer Graphics*, 22(4):205–212, 1988.
7. L. Piegl and W. Tiller. *The NURBS Book*. Springer Verlag, 1997.
8. G. K. Rohde, A. Aldroubi, and B. M. Dawant. Adaptive free-form deformations for inter-patient medical image registration. In *Proc. Medical Imaging: Image Processing*. SPIE, 2001. In press.
9. C. R. Maurer Jr., D. L. G. Hill, A. J. Martin, H. Liu, M. McCue, D. Rueckert, D. Lloret, W. A. Hall, R. E. Maxwell, D. J. Hawkes, and C. L. Truwit. Investigation of intraoperative brain deformation using a 1.5T interventional MR system: preliminary results. *IEEE Transactions on Medical Imaging*, 17(5):817–825, 1998.
10. J. A. Schnabel, C. Tanner, A. Castellano Smith, M. O. Leach, C. Hayes, A. Degenhard, R. Hose, D. L. G. Hill, and D. J. Hawkes. Validation of non-rigid registration using Finite Element Methods. In *Proc. Information Processing in Medical Imaging (IPMI'01)*, volume 2082 of *Lecture Notes in Computer Science*, pages 344–357. Springer Verlag, 2001.