Non-Parametric Discrete Registration with Convex Optimisation

Mattias P. Heinrich[1], Bartlomiej W. Papież[2], Julia A. Schnabel[2], and Heinz Handels[1]

[1] Institute for Medical Informatics, University of Luebeck, Germany
[2] Institute of Biomedical Engineering,
Department of Engineering, University of Oxford, UK
heinrich@imi.uni-luebeck.de,www.mpheinrich.de

Abstract. Deformable image registration is an important step in medical image analysis. It enables an automatic labelling of anatomical structures using atlas-based segmentation, motion compensation and multimodal fusion. The use of discrete optimisation approaches has recently attracted a lot attention for mainly two reasons. First, they are able to find an approximate global optimum of the registration cost function and can avoid false local optima. Second, they do not require a derivative of the similarity metric, which increases their flexibility. However, the necessary quantisation of the deformation space causes a very large number of degrees of freedom with a high computational complexity. To deal with this, previous work has focussed on parametric transformation models. In this work, we present an efficient non-parametric discrete registration method using a filter-based similarity cost aggregation and a decomposition of similarity and regularisation term into two convex optimisation steps. This approach enables non-parametric registration with billions of degrees of freedom with computation times of less than a minute. We apply our method to two different common medical image registration tasks, intra-patient 4D-CT lung motion estimation and inter-subject MRI brain registration for segmentation propagation. We show improvements on current state-of-the-art performance both in terms of accuracy and computation time.

1 Introduction and Background

Deformable image registration is an integral part of medical image analysis and has been the focus of a large amount of research. Most deformable registration algorithms consist of three parts: similarity metric, optimisation strategy and transformation model, a comprehensive overview of the current literature is given in [21]. Image registration between an target image I_t and moving image I_m can in general be stated as an energy optimisation problem. The spatial transformation $\phi = Id + \mathbf{u}$, consisting of identity transform Id and deformation field $\mathbf{u}$, is sought that minimises the cost function $E(I_t, I_m, \mathbf{u})$.

We define a cost function for deformable registration, which consists of a similarity term $\mathcal{S}$ and a regularisation term $\mathcal{R}$, where α is a positive factor that

balances the weighting between both penalties:

$$E(\mathbf{u}) = \mathcal{S}(I_t, I_m, \mathbf{u}) + \alpha\mathcal{R}(\mathbf{u}) \tag{1}$$

Most registration approaches optimise this combined cost function directly. A pointwise similarity term, such as the sum of squared differences, defined for each voxel $\mathbf{x}$: $\mathcal{S}(\mathbf{x}) = (I_t(\mathbf{x}) - I_m(\mathbf{x}+\mathbf{u}))^2$, contains a non-linearity with respect to $\mathbf{u}$. Therefore, it cannot be directly solved using convex optimisation.

Continuous optimisation: One approach is to linearise the similarity term to obtain an approximate step towards the desired solution. This linearisation is only valid for a small deformation step, the registration thus requires a series of small iterative updates. This results in two disadvantages: large deformations of small anatomical features might be lost, as only a local optimum is found, and many iterations are required to reach convergence. Additionally the linearisation requires the derivative of the similarity to be computed, which restricts the flexibility of continuous optimisation approaches.

MRF-based optimisation: In order to overcome the above discussed disadvantages of gradient-based techniques, the use of discrete optimisation has been proposed based on a Markov random field (MRF) formulation [7]. Here, the deformation field $\mathbf{u}$ is not represented by a continuous vector field, but defined as a (dense) set of discrete spatial displacement labels $\mathbf{d} \in \mathcal{L} = \{0, \pm 1, \dots \pm l_{max}\}$[3]. The cost of assigning a certain label $\mathbf{d} \in \mathcal{L}$ to each voxel $\mathbf{x}$ depends on the unary potentials, which correspond to the (pointwise) image similarity and pairwise (or higher-order) interactions (representing the regularisation term). Different optimisation methods can be employed, however, even the simplest inference algorithms (e.g. dynamic programming on a tree [9]) have a high computational complexity. Therefore most previous approaches for discrete optimisation in medical image registration used parametric transformation models[3]. Glocker et al. [7] used a B-spline transformation model and further reduced the complexity by restricting the potential displacements to lie along the three normal axes. In [17] a finite element method is used to parametrise the transformations.

Local cost aggregation: Alternatively, local regularisation models have been used for discrete registration (or stereo estimation). In cost aggregation approaches [19], no explicit global regularisation of the deformation field is performed, and the motion is only locally constrained to be smooth by averaging the similarity term over a small window. These techniques work reasonably well and are computationally fast. However, they often require ad-hoc post-processing steps to propagate information into textureless regions.

In this work, we present a different solution, which is based on a dual optimisation technique following the idea of [5]. The non-convex optimisation problem can be split into two convex problems: one for the similarity term and one for the regularisation, by introducing an auxiliary vector field linking both terms. Each term is optimised in alternation, increasing the linking between them until both vector fields converge. This concept has been previously investigated in [2]

[3] The only exception for 3D registration we know of is [20], where a non-parametric solution is obtained with graph cuts resulting in computation times of 24 hours.

resulting in the "Pair-And-Smooth, Hybrid Energy based Algorithm" (PASHA), and can be implicitly found in the demons registration approach [24]. However, these approaches are based on gradient-based optimisation of the similarity cost. In contrast, we propose to use a discretised search space to find a better optimum of the similarity term (following the ideas presented in [22] for optical flow estimation) and additionally use a local cost aggregation step.

The remainder of this paper is organised as follows. Sec. 2 introduces our method consisting of local cost aggregation, convex optimisation for global smoothness and minimisation of inverse inconsistencies. The implementation used to efficiently solve the proposed model will be explained in detail. In Sec. 3 we present the employed experiments on two different challenging medical image registration tasks. Sec. 4 discusses the results and gives an outlook on further possible research directives.

2 Methods

Our proposed non-parametric discrete registration method aims to find the deformation field $\mathbf{u}$, which minimises a cost function over the image domain:

$$E(\mathbf{u}) = \sum_{\Omega} \mathcal{S}(I_t, I_m, \mathbf{u}) + \alpha |\nabla \mathbf{u}|^2 \tag{2}$$

The second term penalises the squared gradient of the displacement field, which forms a diffusive regularisation. The similarity term can be either point-wise or defined over a local neighbourhood (image patch) Ω.

Our method estimates a deformation field in three steps. First, an explicit search is performed over a discrete displacement space only enforcing smoothness by **local cost aggregation** (Sec. 2.1). This yields a prior map of probable displacements based on the similarity term and an initial best displacement for each voxel. Second, **global smoothness** is enforced through alternative updates of the estimated deformations and the similarity distribution using an auxiliary term (Sec. 2.2). Third, **inverse consistency** [6] is achieved by minimising the discrepancy between forward and backward transforms (Sec. 2.3). This can be used to enforce a one-to-one mapping between two images, in order to avoid physically implausible folding of the deformation field.

2.1 Local cost aggregation

Following the definition of previous discrete registration approaches, e.g. [9], we restrict the deformations $\mathbf{u}$ to be part a quantised set of 3D displacements $\mathbf{d} \in \mathcal{L}$ for each voxel $\mathbf{x}$:

$$\mathbf{d} \in \mathcal{L} = \{0, \pm q, \pm 2q, \ldots, \pm l_{max}\}^3,$$

with a quantisation step q and maximal displacement range of l_{max}. The advantage of this approach compared to a linearisation of the similarity function is that both an iterative solution and a local optimum can be avoided. Analogously,

to previous work on local cost aggregation in stereo estimation [12, 19], we first construct a six-dimensional displacement space volume (DSV), which extends the image dimensions by three dimensions (for the 3D displacement label space $\mathcal{L}$). Each entry of the DSV represents the point-wise similarity cost of translating a voxel $\mathbf{x}$ with a certain displacement $\mathbf{d}$:

$$DSV(\mathbf{x}, \mathbf{d}) = \mathcal{S}(I_t(\mathbf{x}), I_m(\mathbf{x} + \mathbf{d})) \tag{3}$$

To enforce constant motion within a small local region, we average the similarity term over a local patch for every voxel. However, a naive implementation of such a windowed cost evaluation would have considerable computational cost. Yet, for uniform or spatially weighted (e.g. B-Spline) patches a cost aggregation with constant complexity regardless of the patch size can be obtained through a spatial convolution (or moving average) filter. The convolution filter K is applied to every 3D subvolume (i.e. in spatial domain) of the DSV with a constant displacement (see [12], Fig. 1 for a visual example of this procedure).

We could stop here and directly obtain a displacement field by selecting the displacement $\mathbf{d}$ with the lowest aggregated cost for each voxel:

$$\mathbf{u} = \underset{\mathbf{d} \in \mathcal{L}}{\operatorname{argmin}}(K \star DSV(\mathbf{d})) \tag{4}$$

This concept, which is often called winner-takes-all (WTA), already achieves a relatively robust estimation of deformations. An alternative way to obtain a cost aggregation is though iterative diffusion of the similarity images [19], which repeatedly replaces the value of a voxel with a weighted average of its neighbours.

The concept of filter-based local cost aggregation is directly suitable for point-wise similarity metrics. One possible choice is to use the sum of absolute differences (SAD) of self-similarity context (SSC) descriptors as introduced by Heinrich et al. [8]. The descriptors represent the self-similarity of small image patches within a local neighbourhood of each voxel. This yields a twelve-valued vector, which is quantised into a single 64-bit integer. Evaluating the SAD of two vectors $\mathbf{s}_t$ and $\mathbf{s}_m$ with quantised representations S_t and S_m therefore simplifies to calculating their Hamming weight:

$$SSC(\mathbf{x}) = \sum_{i=1}^{12} |s_t(\mathbf{x})^i - s_m(\mathbf{x})^i| = \Xi\{S_t(\mathbf{x}) \oplus S_m(\mathbf{x})\}, \tag{5}$$

where the function $\Xi\{\cdot\}$ represents the bit-count operation. The advantage of this metric is that it can be evaluated very quickly and has been shown to be robust against local change in contrast and image noise. A local cost aggregation within the DSV follows accordingly to Eqs. 3 and 4.

We can extend the local cost aggregation from point-wise to patch-based similarity metrics. In particular local cross-correlation (LCC) is here of interest, as it has been widely used for medical image registration [1, 15]. LCC can be directly calculated over a local window Ω centred around $\mathbf{x}$ by:

$$LCC(\mathbf{x}) = \frac{\sum_\Omega (I_t(\mathbf{x}) - \mu_t)(I_m(\mathbf{x}) - \mu_m)}{\sqrt{\sum_\Omega (I_t(\mathbf{x}) - \mu_t)^2} \sqrt{\sum_\Omega (I_m(\mathbf{x}) - \mu_m)^2}} \tag{6}$$

where μ_m and μ_t define the local intensity mean in moving and target images respectively. These values and the local standard deviations $V_{m,t}(\mathbf{x}) = \sqrt{1/|\Omega| \sum_\Omega (I_{m,t}(\mathbf{x} - \mu_{m,t})^2}$ do not depend on the translational displacement $\mathbf{d}$ and can therefore be pre-computed once for both images. Following the approach in [14] (similarly presented in [1]), we can efficiently compute LCC for each displacement in constant complexity independent of the patch-size. When expanding the numerator of Eq. 6, we obtain:

$$\frac{1}{|\Omega|} \sum_\Omega I_t(\mathbf{x}) I_m(\mathbf{x}) - \mu_t \sum_\Omega I_m(\mathbf{x}) - \mu_m \sum_\Omega I_t(\mathbf{x}) + \mu_t \mu_m.$$

Since $\mu_{t,m} = 1/|\Omega| \sum_\Omega I_{t,m}(\mathbf{x})$, we can simplify this to: $1/|\Omega| \sum_\Omega I_t(\mathbf{x}) I_m(\mathbf{x}) - \mu_t \mu_m$. The summation of the first term, can again be more efficiently computed by first taking the point-wise product of the image intensities followed by a constant time averaging filter. Since, the local variances have been pre-computed, evaluating the LCC for each voxel and displacement only requires 10 operations: a huge speed-up compared to the naive approach, especially for large windows.

Using the WTA approach, however, does not enforce any global smoothness and can therefore lead to poor motion estimation for homogenous areas with little texture. For these reasons, most local stereo estimation methods perform post-processing steps to remove false correspondences.

2.2 Global smoothness with convex optimisation

To improve the motion estimation for homogenous areas, we adopt the approach presented by [22], which follows the primal dual approaches for total variation based image processing [5]. An auxiliary second deformation field $\mathbf{v}$ is introduced and the combined cost function $E(\mathbf{v}, \mathbf{u})$ is solved in two alternating steps.

$$E(\mathbf{v}, \mathbf{u}) = DSV(\mathbf{v}) + \frac{1}{2\theta}(\mathbf{v} - \mathbf{u})^2 + \alpha |\nabla \mathbf{u}|^2 \tag{7}$$

The optimal selection of $\mathbf{v}$ with respect to the similarity term, together with the auxiliary middle term can be performed globally optimal, as before, using local cost aggregation and WTA selection (of the DSV plus the coupling term). Note, that the disparity space volume (DSV) has to be computed only once. The regularisation penalty can be solved optimally by a Gaussian smoothing of the deformation field. The parameter α controls the diffusivity of the deformation field and is implicitly set through the variance of the Gaussian kernel σ^2. The update is performed by $\mathbf{u} \leftarrow K_\sigma \mathbf{v}$. The parameter θ models the coupling between similarity and regularisation penalty and is decreased during a number of iterations. In our experiments, we have used five iterations of this dual convex optimisation with $\theta = \theta_0 \cdot \{150, 50, 15, 5, 1.5\}$, where θ_0 is a parameter which should be adapted to the range of a specific similarity metric. For $\theta \to 0$ we reach convergence and $\mathbf{u} = \mathbf{v}$. We have chosen isotropic diffusion regularisation with Gaussian smoothing, which has been widely used for medical image registration, but other regularisation penalties or filters (e.g. bilateral filters [16]) could be easily integrated into this framework.

2.3 Inverse consistency

Following our above formulation, the registration outcome would be dependent on the choice of target and moving image. To remove this bias and ensure a one-to-one mapping, we use a simple scheme to obtain inverse consistent mappings, given the forward and backward displacement fields $\mathbf{u}_f$ and $\mathbf{u}_b$ respectively (which are independently calculated). We aim to reduce the inverse consistency error (ICE) [6]. This can be achieved according to [8] by iteratively updating the following equations:

$$\begin{aligned} \mathbf{u}_f^{n+1} &= 0.5(\mathbf{u}_f^n - \mathbf{u}_b^n(\mathbf{x} + \mathbf{u}_f^n)) \\ \mathbf{u}_b^{n+1} &= 0.5(\mathbf{u}_b^n - \mathbf{u}_f^n(\mathbf{x} + \mathbf{u}_b^n)) \end{aligned} \tag{8}$$

where the initial (asymmetric) transformation are denoted by a time-point $n = 0$. Empirically, we found that 10 iterations are sufficient to reduce the ICE to insignificantly low values and also ensure the absence negative Jacobian values (and thus singularities) within the deformation fields. Further details on the convergence of this scheme can be found in [10] Ch. 4.4.1.

3 Experiments

In order to show the benefits of our new approach, we compare its performance for two challenging datasets for medical image registration. First, the deformable registration of inter-patient brain MRI scans, and second, intra-patient lung motion estimation of 4D-CT scans. The first experiment employs the Columbia University Medical Center (CUMC) dataset [4], consisting of 1.5 T MRI scans of 12 subjects, which have been manually labeled into 130 anatomical or functional regions. For the second experiment, we use the ten cases of the DIR-lab dataset [3], which has been validated with 300 manual landmarks for both the maximum inspiration and maximum expiration phase of a breathing cycle. The motivation for choosing these two datasets is that both have been widely used to evaluate state-of-the-art deformable registration techniques.

3.1 Parameter choices

We use a multi-resolution scheme with three levels and downsampling factors of $\{3, 2, 1\}$ for all experiments. We use a **de**ns**e** **d**isplacement **s**ampling (**deeds**) [9] with a sampling range of $l_{max} = \{6, 4, 2\}$ voxels for the three resolution levels and a quantisation of $q = 1$ voxel. For the brain and lung dataset, we chose LCC and SSC [8] respectively as similarity metric. To obtain SSC descriptors in a lower resolution, we calculate self-similarity distances in the original image resolution and downsample only the final descriptors. A parameter variation has been performed for a subset of the CUMC12 registration experiment, to examine the influence of the radius r of the box-filter for local cost aggregation (or patch-size of LCC computation respectively), the parameter σ for the diffusive

regularisation of deformation fields and θ_0 to scale the range of the similarity metric. The Dice metric $D = 2|A \cap M|/(|A| + |M|)$ between automatic and manual segmentations A and M (calculated separately for each of the 130 label regions) has been chosen to evaluate the quality of registrations. The registration accuracy varies between $D = 50.5 - 50.8\%$ for $0.4 \leq \sigma \leq 1.0$ voxels and $D = 50.4 - 50.8\%$ for $0.25 \leq \theta \leq 4$. Using radii of $\{0, 1, 2, 3, 4\}$ voxels for the LCC metric results in segmentation overlaps of $\{36.0, 50.9, 50.8, 50.3, 49.3\}\%$, indicating a very good robustness except when skipping the cost aggregation entirely. The chosen parameters for all further experiments using the symmetric formulation are: $\sigma = 0.6$, $r = 2$, and $\theta_0 = 1$. The best settings for an asymmetric registration are the same except for $\sigma = 1.2$, which results in $D = 50.8\%$.

3.2 Computation time

We use an efficient multi-threaded CPU implementation, which is being made available at www.mpheinrich.de/software.html and run experiments on a dual-core processor. When using an asymmetric registration formulation, $\approx$ 20 sec. (for all resolutions) are spent on calculating the DSV and aggregating its cost locally, and half a minute for the iterative global regularisation, yielding a total time for one 3D brain registration (with a volume size of $256 \times 256 \times 124$) of less than one minute. If a symmetric transformation is required, these computation times double and an additional time of $\approx$ 40 sec. has to be added for enforcing inverse consistency of the deformation fields. When using the SSC similarity metric, $\approx$ 10 sec. are spent on calculating the descriptors. These processing times are more then an order of magnitude faster than the top performing algorithms in [13] and further speed-ups could be expected when using a GPU (c.f. [12]).

3.3 Results

The CUMC12 dataset has been used in a comprehensive comparison study of 14 non-linear registration methods in [13]. In those experiments, the Jaccard index $J = (A \cap M)/(A \cup M)$ was used to evaluate the registration accuracy. We are therefore able to directly compare our approach to 14 other algorithms for a total of 132 one-to-one registrations. The same pre-processing of the data as detailed in [13] has been used: in particular removing the skull in the images using the provided brain masks. Our method achieves the highest overlap of all methods, with $J = 36.3\%$, see Fig. 1. The usefulness of the global regularisation step is demonstrated by performing the same experiments as before, but this time only use the local cost aggregation (see Eq. 4) followed by one Gaussian smoothing as post-processing. This variant (denoted as "local") achieves only $J = 34.5\%$. We also outperform the FEM based discrete registration approach of, Popuri et al. [17], whose work is the most similar to ours and achieves $J = 31.4\%$. In contrast to us, they use a parametric transformation model, variational smoothing and do not include a global convex optimisation of the regularisation term.

The second experiment is challenging for continuous optimisation approaches, because there is a large discrepancy of the magnitude and direction of the motion

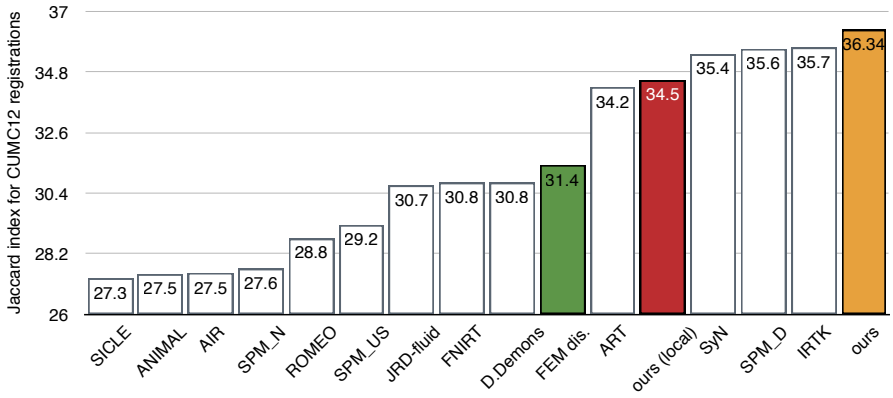

Fig. 1: Jaccard overlap (over 130 regions) for 16 non-linear registration algorithms for CUMC12 dataset. Our approach achieves the highest accuracy with $J = 36.3\%$.

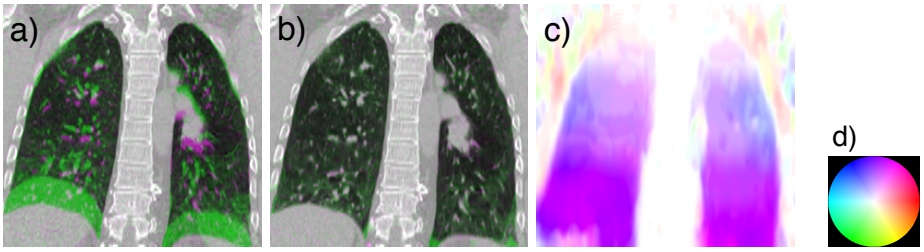

Fig. 2: Coronal view of overlay of inhale and exhale phase of Case # 6 of [3] before (a) and after (b) alignment using our proposed non-parametric discrete registration. The estimated deformation field (c) is represented by HSV colours (d), where the vector orientation is indicated by hue and the deformation length by saturation. The complexity of the deformation field std(Jac) lies between 0.23 and 0.33 for the 10 cases.

inside and outside of the lungs. Currently, most approaches that achieve a high accuracy, e.g. [11] or [18] deal with this problem by segmenting the lungs and masking out the rib-cage and other body parts. In our approach, a discrete sampling of the displacement, with a very large range of possible motion vectors $(6 \times 2 + 1)^3 = 2197$, is used in the lowest resolution to capture large motion of small features. Using the same settings as before (except that we selected SSC, as we found it works better as similarity metric for this task), we obtain a target registration error of 1.17 mm. This is only marginally higher than the best results, previously achieved for masked registration: 0.99 mm ([11] and [18]). Our approach is so far the most accurate for unmasked registration for this dataset, with an improvement of $\approx$ 0.25 mm to the results of [9] (1.43 mm) and [11] (1.41 mm w/o masks). Figure 2 shows an example registration outcome including the obtained deformation field.

4 Conclusion

We have presented a new framework for discrete medical image registration, which includes both local and global regularisation. The search space of potential deformations is sampled in a dense manner, thus avoiding local minima or the need for an iterative refinement. A local regularisation is integrated by a cost aggregation scheme, which is performed through a spatial filtering of the displacement space volume (DSV). The global diffusive smoothness prior is enforced through an alternating update of the distribution of the locally aggregated image similarity and a global deformation field smoothing through Gaussian convolution. Solving each of the two decoupled functionals separately results in convex optimisation problems that can be solved optimally. After few iterations, this procedure converges to a very good approximation of the optimum of the combined cost function (and a substantial improvement over using only local regularisation). Our approach, which achieves computation times of less than one minute per 3D registration, performs best on the CUMC12 brain dataset in comparison to 15 other state-of-the-art techniques and within 0.2 mm of the best approaches for the DIR-Lab lung dataset. Additional results for the remaining three datasets of the Klein study [13] support the initial findings resulting in Jaccard scores of 56.38% (for LPBA40), 39.54% (for MGH10) and 36.43% (for IBSR18), which each outperform the previously best results.

Further research could improve on the presented results, by integrating additional information, e.g. segmentations or better priors on the deformation field regularity. In the future, we would like to directly compare our approach to global MRF-based optimisation strategies [7, 9]. The use of this framework for other challenging medical image registration tasks, including multi-modal registration, is be directly possible. An interesting alternative to the use of identical support regions for each voxel could be the use of multiple potential window sizes to represent simultaneously multiple scales of deformations (c.f. [23]).

Acknowledgements. B.W.P. and J.A.S. would like to acknowledge funding from the CRUK/ EPSRC Cancer Imaging Centre at Oxford.

References

1. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ANTs similarity metric performance in brain image registration Neuroimage 54(3), pp. 2033–2044 (2011)
2. Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N.: Iconic feature based nonrigid registration: The PASHA algorithm Comput. Vis. Image Underst. 89 (2-3), pp. 272–298 (2003)
3. Castillo, R., Castillo, E., Guerra, R., Johnson, V.E., McPhail, T., Garg, A.K., Guerrero, T.: A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. Phys. Med. Biol. 54(7), 1849 (2009)
4. Caviness, Jr V.S., Meyer, J., Makris, N., Kennedy, D.N.: MRI-based Topographic Parcellation of Human Neocortex: An Anatomically Specified Method with Estimate of Reliability. Journal of Cognitive Neuroscience 8(6), pp. 566–587 (1996)

5. Chambolle, A.: An algorithm for total variation minimization and applications. Journal of Mathematical Imaging and Vision 20(1-2), pp. 89–97 (2004)
6. Christensen, G.E., Johnson, H.J.: Consistent Image Registration IEEE Trans. Med. Imag. 20(7), pp. 568–582 (2001)
7. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. Med. Imag. Anal. 12(6), pp. 731–741 (2008)
8. Heinrich, M.P., Jenkinson, M., Papiez, B.W. and Brady, M., Schnabel, J.A.: Towards Realtime Multimodal Fusion for Image-Guided Interventions using Self-Similarities In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. MICCAI 2013, LNCS, vol. 8149, pp. 187–194, (2013)
9. Heinrich, M.P., Jenkinson, M. Brady, M., Schnabel, J.A.: MRF-based Deformable Registration and Ventilation Estimation of Lung CT. IEEE Trans. Med. Imag. 32(7), pp. 1239–1248 (2013)
10. Heinrich, M.P.: Deformable lung registration for pulmonary image analysis of MRI and CT scans. University of Oxford (2013)
11. Hermann, S., Werner, R.: High Accuracy Optical Flow for 3D Medical Image Registration using the Census Cost Function PSIVT 2013, pp. 1–13 (2013)
12. Hosni, A., Rhemann, C., Bleyer, M., Rother, C., Gelautz, M.: Fast Cost-Volume Filtering for Visual Correspondence and Beyond IEEE Trans. Pattern Anal. Mach. Intell., 35(2) pp. 504–511 (2013)
13. Klein, A. et al: Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. Neuroimage 46(3), 786–802 (2008)
14. Lewis, J.P.: Fast normalized cross-correlation Vision Interface 10(1), pp. 120–123 (1995)
15. Lorenzi, M., Ayache, N., Frisoni, G.B., Pennec, X.: LCC-Demons: a robust and accurate diffeomorphic registration algorithm. NeuroImage 81, pp. 470–483 (2013)
16. Papiez, B.W., Heinrich, M.P., Risser, L., Schnabel, J.A.: Complex Lung Motion Estimation via Adaptive Bilateral Filtering of the Deformation Field In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. MICCAI 2013, LNCS, vol. 8151, pp. 25–32, (2013)
17. Popuri, K., Cobzas, D., Jägersand, M. A Variational Formulation for Discrete Registration In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. MICCAI 2013, LNCS, vol. 8151, pp. 187–194, (2013)
18. Rühaak, J., Heldmann, S., Kipshagen, T., Fischer, B.: Highly Accurate Fast Lung CT Registration In: Ourselin, S., Haynor, D.R. SPIE Medical Imaging, pp. 1–9, (2013)
19. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms Int. J. Comput .Vision, 47(1), pp. 7–42 (2002)
20. So, R.W.K, Tang, T.W.H, Chung, A.C.S.: Non-rigid image registration of brain magnetic resonance images using graph-cuts. Pattern Recognition 44(10-11), pp. 2450–2467 (2011)
21. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable Medical Image Registration: A Survey IEEE Trans. Med. Imag. 32(7), 1153–1190 (2013)
22. Steinbrücker, F., Pock, T., Cremers, D.: Large displacement optical flow computation without warping ICCV 2009, pp. 1609–1614 (2009)
23. Veksler, O.: Fast Variable Window for Stereo Correspondence using Integral Images CVPR 2003, pp. 1–6 (2003)
24. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration, NeuroImage, 45(1), pp. 61–72 (2009)