

Multimodal Surface Matching: Fast and Generalisable Cortical Registration Using Discrete Optimisation

Emma C. Robinson¹, Saad Jbabdi¹, Jesper Andersson¹, Stephen Smith¹,
Matthew F. Glasser², David C. Van Essen², Greg Burgess²,
Michael P. Harms³, Deanna M. Barch⁴, and Mark Jenkinson¹

¹ FMRIB, Nuffield Department of Clinical Neurosciences, University of Oxford, UK

² Department of Anatomy and Neurobiology, Washington University School of Medicine, St Louis, MO, USA

³ Department of Psychiatry, Washington University School of Medicine, St Louis, MO, USA

⁴ Department of Psychology, Washington University, St Louis, MO, USA

Abstract. Group neuroimaging studies of the cerebral cortex benefit from accurate, surface-based, cross-subject alignment for investigating brain architecture, function and connectivity. There is an increasing amount of high quality data available. However, establishing how different modalities correlate across groups remains an open research question. One reason for this is that the current methods for registration, based on cortical folding, provide sub-optimal alignment of some functional sub-regions of the brain. A more flexible framework is needed that will allow robust alignment of multiple modalities. We adapt the Fast Primal-Dual (Fast-PD) approach for discrete Markov Random Field (MRF) optimisation to spherical registration by reframing the deformation labels as a discrete set of rotations and propose a novel regularisation term, derived from the geodesic distance between rotation matrices. This formulation allows significant flexibility in the choice of similarity metric. To this end we propose a new multivariate cost function based on the discretisation of a graph-based mutual information measure. Results are presented for alignment driven by scalar metrics of curvature and myelination, and multivariate features derived from functional task performance. These experiments demonstrate the potential of this approach for improving the integration of complementary brain data sets in the future.

1 Introduction

Automated and accurate registration of the cortical sheet is an increasingly important topic of neuroimaging methods research. However, the use of volumetric registration approaches in functional studies is sub-optimal, since functional areas are often spaced much further apart across the two-dimensional surface than is represented in volumetric space, because of the cortical folds.

Spherical registration algorithms [1,2] simplify the cortical matching problem by inflating the surface to a two dimensional sphere. Current approaches predominantly use measures of sulcal depth or mean curvature (folding) to perform

the cross-surface matching. This leads to a reasonable first approximation of brain alignment, but is limited by the fact that significant variability in cortical folding manifests across populations [3].

A more fundamental caveat of morphologically driven alignment is that cortical folding imperfectly reflects functional sub-divisions[3]. The fundus of a sulcus does not generally align with the boundary between two functional regions. Since the ultimate goal of inter-subject registration is to match functional regions across subjects, matching anatomies yields only an approximation.

Connectivity-based alignment [4] has been proposed as an additional means for driving registration. Connections may be a more direct correlate of brain function than local morphology. However, it is unclear how well we can estimate connections that are relevant to drive registration. Another alternative is to use functional (task) activation data to drive registration, though at the cost of increasing data requirements and incomplete brain coverage.

It is unlikely that any single measure of brain structure or function will be sufficient for consistent, whole brain, and functionally accurate inter-subject registration. A multimodal approach is likely to yield significant improvements, but requires a more flexible registration framework that is adaptable to multivariate correlates of functional, structural and connectional brain data.

Discrete optimisation approaches to image registration constitute such a flexible framework. They were first proposed by Glocker et al [5] who converted a B-spline free-form deformation model [6] into a discrete setting. Discrete sets of deformation labels were assigned to the control points of the B-spline model, and the problem was re-formulated using graph cuts [7]. Alternative approaches that do not tie the framework to use of the free-form deformation model have also been proposed by [8,9].

Discrete methods in general offer advantages in terms of reduced sensitivity to local minima combined with a fast and efficient optimisation approach that provides flexibility by not restricting the choice of similarity measure. Here, we introduce a multivariate mutual-information measure derived from entropic graphs [10,11]. These measures were first proposed for use in volumetric registration by [10] in order to allow fast estimation of multivariate mutual information without the need for costly estimations of high dimensional histograms. To our knowledge, ours is the first application of multivariate matching within the discrete optimisation framework.

One restriction of the MRF-based optimisation framework is that the regularisation term must be reformulated in terms of pair-wise edge potentials between neighbouring vertices. This can limit the choice of regulariser to measures derived from the first order derivatives of the deformation warp, which are not invariant to scaling or rotations. Alternative methods have been proposed in [12,13]. Here we propose a new regularisation potential based on the geodesic distance between rotation matrices. This is not a metric and therefore we benefit from applying Fast-PD optimisation of the MRF, proposed in [14,15] and first used for registration by Glocker et al in [5].

In the next section we provide an overview of the new Multimodal Surface Matching (MSM) approach and its reformulation to fit the spherical coordinate system. Results are presented using simulated data, univariate measures of sulcal depth and curvature, and finally multivariate features derived from functional task performance.

2 MRF-Based Optimisation for Spherical Registration

2.1 Discrete Optimisation and Fast-PD

Discrete registration methods redefine the image as a weighted graph $G(P, E, \omega)$ formed from control point nodes (P), edges (E), and weights (ω), and solve the optimisation in the form of a Markov Random Field labelling. This approach limits the movement of any vertex point p to a set of discrete displacements $\mathbf{d}^{l_p}$ determined by a predefined label set $L = \{a, b, \dots\}$. For spherical registration these displacements may be governed by a finite set of rotations (Fig . 1).

The optimal deformation can be found by minimising the following energy function:

$$COST = \sum_{p \in P} \mathbf{c}_p(l_p) + \sum_{(p,q) \in E} \omega_{pq} V(l_p, l_q) \quad (1)$$

This is formed from a similarity term $\mathbf{c}_p(l_p)$ and a weighted ω_{pq} penalisation term $V(l_p, l_q)$, which limits the extent to which neighbouring vertices (p, q) can be assigned different displacement labels (l_p, l_q). The penalisation term controls regularisation of the deformation. The key advantage of the MRF formulation is that any similarity metric may be used.

Volumetric methods, such as [5,8] propose using a low-resolution control point grids or regular sampling of the image space to define the nodes in G . In this way image similarity, for each label a , is approximated by factorizing the expression for whole image similarity into cliques, as:

$$\sum_{p \in P} \mathbf{c}_p(l_p) = \sum_{p \in P} \sum_{i \in N} \hat{n}(\mathbf{x}_i - \mathbf{x}_p) (sim(F(\mathbf{x}_i), M((\mathbf{x}_i + \mathbf{d}^a))) \quad (2)$$

This sums over pair-wise similarities between points in the moving mesh $M((\mathbf{x}_i + \mathbf{d}^a))$ (transformed by the proposed deformation $\mathbf{d}^a$) and their neighbours within the fixed/target image $F(\mathbf{x}_i)$, for all data points in the neighbourhood of control point p . The $\hat{n}(\mathbf{x}_i - \mathbf{x}_p)$ is a weighting term controlling the zone of influence of each control point. In this instance, $\hat{n}(\mathbf{x}_i - \mathbf{x}_p)$ is either 0 or 1.

There are several methods for solving the MRF problem, amongst which are the well-known graph cut [7] and max-flow min-cut algorithms, used often in computer vision for image restoration and segmentation. In [5] the authors use the Fast-PD algorithm, proposed by Komodakis et al [14,15] which takes advantage of the primal-dual schema of linear programming to derive an efficient approximation. For full comprehensive details of the algorithm we refer to [14,15]. However, the main advantage is that it generates fast solutions

that can be shown to be close to the global optimum even for pair-wise potential terms which are not metrics. That is, potentials can be used that do not have to satisfy the conditions $V(a, b) = 0 \ \forall a = b$; $V(a, b) = V(b, a) \geq 0$ and $V(a, b) = V(b, a) \leq V(a, c) + V(c, b)$.

2.2 Deformation Labels and the Control Point Grid

In this framework, spherical registration is driven using a series multi-resolution control point grids. Currently three regular icosahedrons of order 3-5 (161, 642 and 2542 vertices) are used, with respective mean vertex distances (MVD) of 26.7mm, 13.8mm and 6.9mm. Similarity is estimated by dividing all vertices of the original mesh into sets of cliques as for eq. 2. Interpolation of the control point warp to the original surface mesh is performed using spline-based interpolation.

Image data is downsampled to a 10,000 vertex regular mesh ($MVD = 3.5\text{mm}$) to speed up the sampling of the image space. In general, this is a reasonable approximation as the original volumetric image data will typically be of a resolution between 1mm for structural data and 3mm for functional data. Downsampling is performed using Gaussian interpolation, which enables simultaneous smoothing of the data. Note, although higher resolution control point and image grids can be used with minimal additional computational overhead, these were not found to improve the results of the data in this paper.

Deformation labels are defined via a simple and approximate solution to sampling on the sphere: a higher resolution icosahedron grid is projected below each control point and used to define the sampling grid (see Fig. 1). MRF labels are represented as discrete rotations between the control point and the vertices of the higher resolution sampling grid. For a sampling grid formed from an icosahedron 2 orders higher than the control point grid, there is typically between 10 and 20 labels per control point. The spacing between labels at the highest resolution level is roughly 1mm.

At each level the maximum sampling distances, and thus biggest possible control point deformation, is set to $0.4 * MVD$. This ensures diffeomorphic warps by forcibly preventing mesh folding. However, this also means that it is not guaranteed that optimal alignment can be reached within one cycle of labelling. Therefore, the algorithm iterates over several cycles of labelling at each resolution. As the MRF will be optimised for the label set after each cycle, the label set iterates between using the vertices and barycentres of the faces of the sampling grid to prevent the registration getting trapped in local minima. After each stage the deformation is projected to the image grid and the control point grid is reset.

2.3 Regularisation

Regularisation in discrete optimisation methods has conventionally been imposed by pair-wise potentials calculated from the distance between the deformations assigned to neighbouring points: $V(l_p, l_q) = \lambda(|d^{l_p} - d^{l_q}|)$. However, in methods such as this where registration iterates over several cycles of MRF optimisation, symmetric pairwise potentials are unsuited as they cannot capture




Fig. 1. The sampling (blue) and control point grids (black). Sampling points vary between the vertices and the barycentres (red points) of the high res grid. Max sampling distance is 0.4 times the control point spacing. Rotation matrices, estimated between the control point and each of the sampling points, are used to define the deformation labels.

the properties of the previous warps. Estimates of the deformation from previous iterations (R_p, R_q) can be used to provide a full regularisation over time as: $V(l_p, l_q) = \lambda(|\mathbf{R}_p d^{l_p} - \mathbf{R}_q d^{l_q}|)$ [5]. Nevertheless, this automatically renders any distance non-metric as the penalisation need no longer be same for points assigned the same label: $V(l_p, l_q) \neq 0$ if $l_p = l_q$.

In the spherical registration format, it is no longer desirable to project deformations as Cartesian deformation vectors, as the deformation is constrained to the spherical surface, and thus the space is not Euclidean. Instead we choose to represent the deformation labels as rotation matrices defined using the Rodrigues rotation formula:

$$\mathbf{R}_p^{l_p} = I + \sin\theta_{l_p} [\mathbf{k}_{l_p}]_{\times} + (1 - \cos\theta_{l_p})(\mathbf{k}_{l_p} \mathbf{k}_{l_p}^T - I), \quad (3)$$

where, the axis $\mathbf{k}_{l_p} = \mathbf{x}_p \times \mathbf{x}_s$ and angle θ_{l_p} of rotation for label l_p are estimated using the initial $\mathbf{x}_p$ and final point of the rotation $\mathbf{x}_s$ (Fig. 1). $[\mathbf{k}_{l_p}]_{\times}$ is the cross product matrix.

This form is also conveniently linked to the definition of a geodesic distance between rotation matrices: $d_g(\mathbf{A}, \mathbf{B}) := \|\log(\mathbf{A}^T \mathbf{B})\|_F = \|[\mathbf{k}_{ab}]_{\times} \theta_{ab}\|_F$, where $\|\cdot\|_F$ represents the Frobenius norm and $d_g(A, B)$ is proportional to the angle of difference (θ_{ab}) between the rotation matrices. This allows warps from previous iterations to contribute to the estimation of the pairwise potential. As the control point grid is reset each time, the past deformation of each control point p is estimated from the warped imaged data, and summarised by rotation matrix $\mathbf{R}_p$. The full rotation matrix for the combined deformation can then be estimated as $\mathbf{A} = \mathbf{R}_p^{l_p} \mathbf{R}_p$, and compared against the full deformation of its neighbour $\mathbf{B} = \mathbf{R}_q^{l_q} \mathbf{R}_p$.

3 Multivariate Similarity Measures

A significant advantage of the flexibility of the MRF framework is that it is straightforward to replace scalar similarity terms with a multivariate measure.

Therefore, our framework can incorporate a multivariate measure such as the estimate of mutual information known as α -entropy, derived by [10] from entropic graphs. Rényi, or α -entropy is a generalisation of Shannon's entropy of the form: $H_\alpha(z) = \frac{1}{1-\alpha} \log(\sum_{i=1}^n p_i^\alpha)$, where z is a random variable. Entropic graphs are any graph whose normalised total edge weight is a consistent estimator of α -entropy. These include the minimal spanning tree (MST) and the k -nearest neighbour (kNN) graph.

The basic principle is as follows: if each point in the fixed image is described by a multivariate feature vector $\mathbf{z}^f(\mathbf{x}_i) = [z_1^f(\mathbf{x}_i) \dots z_d^f(\mathbf{x}_i)]$ of dimension d , and the equivalent point in the moving image is $\mathbf{z}^m(\mathbf{x}_i + \mathbf{d}^a)$, then a graph-based measure of mutual information may be calculated from the kNN graph as:

$$\alpha - \hat{MI}(Z_f, Z_m, Z_{fm}) = \frac{1}{\alpha - 1} \log \frac{1}{N^\alpha} \sum_{i=1}^{N_f} \left(\frac{\Gamma_i^{fm}}{\sqrt{\Gamma_i^f \Gamma_i^m}} \right)^\gamma \quad (4)$$

For:

$$\Gamma_i^f = \sum_{n=1}^k \|\mathbf{z}^f(\mathbf{x}_i) - \mathbf{z}^f(\mathbf{x}_{in})\| \quad (5)$$

$$\Gamma_i^m = \sum_{n=1}^k \|\mathbf{z}^m(\mathbf{x}_i + \mathbf{d}^a) - \mathbf{z}^m(\mathbf{x}_{in} + \mathbf{d}^a)\| \quad (6)$$

$$\Gamma_i^{fm} = \sum_{n=1}^k \|\mathbf{z}^{fm}(\mathbf{x}_i, \mathbf{x}_i + \mathbf{d}^a) - \mathbf{z}^{fm}(\mathbf{x}_p, \mathbf{x}_{in} + \mathbf{d}^a)\| \quad (7)$$

Here, $\mathbf{z}^{fm}$ is the result of concatenating the feature vectors from the fixed and moving image at the transformed position $(\mathbf{x}_i + \mathbf{d}^a)$. Z_f , Z_m , and Z_{fm} represent the set of feature vectors for each kNN graph, $\gamma = d(1 - \alpha)$ and α are user defined tuning parameters, and Euclidean distances are estimated between feature vectors. Currently each feature is given a equal weighting when estimating the cost. However, given that the end goal is to combine complementary data sets, each of which are shown to be consistent across subjects in different areas, this term could be replaced by a weighted sum of squares.

Approximate k -nearest neighbours are calculated using [16]. This method has been shown to lead to a significant boost in speed in solving for k , provided a small error in the graph estimation is tolerable. For this purpose approximate neighbours are sufficient as the principal goal of using the α -MI measure to drive the registration is to ensure that the spatial distribution of image features is the same for source and target meshes.

4 Validation

In this section we test the algorithm on scalar image data and use a simulation to show the algorithm performs as well as the state of the art in this respect.

Table 1. Comparing the new Multimodal Surface Matching method against the state of the art. Performance is judged in terms of Dice overlap results for simulated data, and run times for real data (averaged over 14 test sessions).

	FreeSurfer	Spherical Demons	MSM
Dice Overlap	N/A	0.892 ± 0.002	0.922 ± 0.002
Speed	> 1hr	2 min	4 min

We test the performance of the algorithm against FreeSurfer [1] and Spherical Demons (SD) [2], measuring success in terms of speed and generalisability to new data sets. Given that curvature metrics have been shown to be highly variable across the population, and in the absence of any better neurological marker of ground truth, we use a simulated test case to quantify alignment. We then extend the comparison to standard measures of sulcal depth and curvature, as well as a new data set composed of estimates of cortical myelination [17] to provide more concrete evidence that the algorithm performs well on real data.

4.1 Simulated Deformations

A test case (Fig. 2 a), with colour patches each representing a different scalar value, was created and then transformed by applying a known deformation, in order to generate a target (Fig. 2 b). The deformation was obtained through registration of real image data to the fsaverage template using the FreeSurfer algorithm. The discrete optimisation approach was run for three resolution levels, with 5 internal iterations at each level. At each internal iteration the algorithm reset the control point grid and recalculated the label set as described in section 2.2. The data cost term was estimated using cross correlation. We found the framework largely insensitive to the choice of regularisation parameter. The results of the registration (fig. 2 c) were compared with those obtained from applying Spherical Demons (fig. 2 d), run using the default parameters. FreeSurfer was not used as it could not be adapted to accept non-curvature data.

The accuracy of each approach was assessed using the mean Dice overlap of the sampling patches (Table 1 top row). In general, although both approaches estimate the affine component of the registration correctly, Spherical Demons fails to capture the majority of the non-linear deformation (fig. 2 d). Some investigation showed this to be linked to the regularisation; controlled by modifying the number of times a smoothing kernel is applied to the deformation warp in the second step of the Demons algorithm. Setting the number of smoothing iterations to one managed to capture some, but not all, of the remaining non-linear warp.

4.2 Experiments on Real Data

For the next experiment we used 10 subjects (four with repeat sessions) acquired as pilot data for the WU-Minn Human Connectome Project (HCP). The HCP (<http://humanconnectome.org>) is collecting a large and comprehensive database




Fig. 2. Simulation. Colours represent scalar values in the range 0-4 with borders showing partial volume values




Fig. 3. Results of univariate registration. Top row: cross-subject mean atlases generated by convexity-driven FreeSurfer alignment. Centre row: averaged results after driving Spherical Demons using (from left to right) convexity, curvature and myelin data. Bottom row: results of registration using MSM. MSM driven myelin alignment pulls out structures (circled) that are not present after convexity-driven FreeSurfer or myelin-driven Spherical Demons alignment. The MT+ region is much stronger after myelin driven MSM (white arrow).

of neuroimaging data, including high resolution diffusion, task and resting state functional MRI. For each subject MRI scans were obtained over 3 sessions. Pilot data used here includes structural MRI and five task fMRI (tfMRI) scans (working memory, motor, biological motion, language, and social cognition), acquired using multiband acquisition at 2mm isotropic.

The bottom row of Table 1 displays the speed of the different algorithms. Unlike MSM and Spherical Demons, which use multi-resolution control-point grids, FreeSurfer estimates a dense displacement field for every vertex in the high resolution mesh ($\sim 160K$ vertices). Thus it is much slower. FreeSurfer is also limited to registration of sulcal-depth measures as this is seen to be less sensitive to inter-subject variability in folds.

The top row of Fig. 3 shows group-average maps of sulcal depth (yellow = gyral, near the exterior of the brain; blue = buried cortex), curvature (yellow = gyral folds; blue = sulcal fundi), and myelin maps (ratio of T1- and T2-weighted scans) after registration to FreeSurfer's fsaverage surface. In contrast to the sulcal depth and curvature maps, which only reflect cortical shape, the myelin maps [19] highlight regions of functional significance, including the primary sensory areas and the motion sensitive area MT+.

Fig. 3 also shows results from univariate registration using Spherical Demons (middle row) and MSM (bottom row; parametrised as for section 4.1) driven by the three different data sets (sulcal depth/curvature/myelin). In each case, registration is initialised by affine alignment of the subjects' native sulcal depth surfaces to the fsaverage template. The results of the non-linear component of the convexity (left column) and curvature (middle column) driven alignment compare well to the FreeSurfer averages. However, myelin-driven MSM produces a much sharper average than can be achieved through sulcal depth-driven FreeSurfer registration, or myelin-driven Spherical Demons alignment.

5 Multimodal Surface Matching

We now test the performance of the algorithm on multivariate data and present results using a novel dataset derived from functional task performance. The




Fig. 4. Group Z statistic images for the working memory task, before and after registration. After MSM, cluster size grows and absolute value of the peak statistic (pink circle) increase from 4.49 to 5.54.

full set of five tasks collected in the pilot HCP protocol was used. tfMRI data were corrected for volumetric image distortions, registered to standard structural space, mapped to the cortical surface, and regularized by surface-constrained smoothing (2mm FWHM). Statistical analysis was performed using the FEAT tool from FSL [18] whereby task time courses were regressed against the vertex-wise time series using a general linear model (GLM). Z statistic images estimated for the parameters of interest of each task (in this instance 37 correlates of mean activity) were concatenated into one feature vector for each vertex. Z values were squared to downweight the contribution of low intensity, noisy variations when estimating of the image similarity.

A registration target was chosen at random from the dataset in the absence of a suitable average target. Registration was initialised using convexity-driven affine alignment and then driven using the task features over 3 resolution levels, again with 5 iterations at each level. At each level the α -entropy measure was estimated using gamma of 1 as recommended in [10]. The k-nearest neighbour graphs were estimated for 10 neighbours at each level and a 10% error in the neighbour calculation was allowed.

Following tfMRI-constrained registration, a group level general linear model (GLM) was used to look for improvements in alignment across the group. Fig. 4 shows good improvements in the size and peak statistics of the clusters for the working memory task after MSM. Fig. 5 shows results for myelin maps (top row) transformed using the warp estimated for the task data. The source myelin map is initially not well aligned to the target myelin map (see white reference contour), but becomes better aligned, especially in occipital cortex (left on the surface), after registration. The transformed tfMRI activation for the social task (Fig. bottom row) is also better aligned in occipital cortex even though this particular tfMRI activation differs markedly between the source and target individuals. Further work will explore the potential for using independent components derived from resting state fMRI analyses.

6 Discussion

The advent of large data collection projects such as the HCP and UK Biobank means that it has never been more important to provide schemes for fast and accurate matching of a wide range of brain imaging data. However, current methods for cortical surface matching are insufficient, as they are limited to matching scalar measures of cortical shape. In this paper we propose a new, fast and highly generalisable surface registration approach based on discrete optimisation.

The use of discrete optimisation offers significant flexibility to choice of similarity measure, and through the fast-PD framework it is possible to incorporate non-metric regularisation penalties. To this end we have proposed a new smoothness penalty for spherical registration based on penalising the geodesic distance between the rotation matrices that define the deformation labels. We find this generates remarkably smooth warps. Nevertheless, regularisation based on first




Fig. 5. The results of transforming myelin data using task-driven registration. White dots show the boundaries of structures in the target myelin map. Bottom row: the social task fMRI features before and after registration.

order smoothness penalties, as used here and in [5] are limited as they penalise linear transformations [13]. Therefore future work will explore the potential for higher order smoothness terms as proposed in [13].

To our knowledge this is the first discrete registration method that has used multivariate features. Our preliminary analysis suggests that alignment can be improved using features from a battery of functional tasks, and that it has potential both for increasing the statistical power of functional experiments and improving alignment of other functionally correlated datasets.

Further improvements should be attainable by harnessing complementary information from multiple modalities. Myelin maps have features which are consistent across subjects in regions that are highly variable in cortical folding. In addition, neuronal fibres provide the structure which underpins functional communication. Thus combined estimates of structural (via diffusion tractography) and functional connectivity (via resting-state fMRI) also have significant potential.

References

1. Fischl, B., Sereno, M.I., Tootell, R.B.H., Dale, A.M.: High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping* 8(4), 272–284 (1999)
2. Yeo, T., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., Golland, P.: Spherical demons: fast diffeomorphic landmark free surface registration. *IEEE Transaction on Medical Imaging* 29(3), 650–668 (2010)

3. Fischl, B., Rajendran, R., Busa, E., Augustinack, J., Hinds, O., Yeo, B.T.T., Mohlberg, H., Amunts, K., Zilles, K.: Cortical folding patterns and predicting cytoarchitecture. *Cerebral Cortex* 18, 1973–1983 (2008)
4. Petrović, A., Smith, S.M., Menke, R.A., Jenkinson, M.: Methods for tractography-driven surface registration of brain structures. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) *MICCAI 2009, Part I. LNCS*, vol. 5761, pp. 705–712. Springer, Heidelberg (2009)
5. Glocker, B., Komodakis, N., Paragios, N., Navab, N.: Dense image registration through MRFs and efficient linear programming. *Medical Image Analysis* 12, 731–741 (2008)
6. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Transaction on Medical Imaging* 18(8), 712–721 (1999)
7. Boykov, Y., Veksler, O., Sabik, R.: Fast approximate energy minimization via graph cuts. *IEEE PAMI* 23(11), 1–17 (2001)
8. Heinrich, M.P., Jenkinson, M., Brady, M., Schnabel, J.A.: Globally optimal deformable registration on a minimum spanning tree using dense displacement sampling. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012, Part III. LNCS*, vol. 7512, pp. 115–122. Springer, Heidelberg (2012)
9. Cobzas, D., Sen, A.: Random walks for deformable image registration. In: Fichtinger, G., Martel, A., Peters, T. (eds.) *MICCAI 2011, Part II. LNCS*, vol. 6892, pp. 557–565. Springer, Heidelberg (2011)
10. Neemuchwala, H.F.: Entropic graphs for image registration. PhD thesis, University of Michigan (2005)
11. Staring, M., van der Heide, U.A., Klein, S., Viergever, M.A., Pluim, J.P.W.: Registration of cervical mri using multifeature mutual information. *IEEE Transactions on Medical Imaging* 28(9), 1412–1421 (2009)
12. Glocker, B., Komodakis, N., Paragios, N., Navab, N.: Approximated curvature penalty in non-rigid registration using pairwise MRFs. In: Bebis, G., et al. (eds.) *ISVC 2009, Part I. LNCS*, vol. 5875, pp. 1101–1109. Springer, Heidelberg (2009)
13. Kwon, D., Lee, K.J., Yun, I.D., Lee, S.U.: Nonrigid image registration using *dynamic* higher-order MRF model. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008, Part I. LNCS*, vol. 5302, pp. 373–386. Springer, Heidelberg (2008)
14. Komodakis, N., Tziritas, G.: Approximate labeling via graph cuts based on linear programming. *PAMI* 29(8) (2007)
15. Komodakis, N., Tziritas, G.: Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal dual strategies. *Computer Vision and Image Understanding* 112(1), 14–29 (2008)
16. Arya, S., Mount, D., Netanyahu, N., Silverman, R., Wu, A.: An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *J. ACM* 45(6), 891–923 (1998)
17. Glasser, M.F., Van Essen, D.C.: Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI. *J. Neuroscience* 31(32), 11597–11616 (2011)
18. Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M.: FSL. *NeuroImage* 62(2), 782–790 (2012)