

Multimodal Surface Matching with Higher-Order Smoothness Constraints[☆]

Emma C. Robinson^{a,g,h,*}, Kara Garcia^{b,1}, Matthew F. Glasser^{e,d,1}, Zhengdao Chen^e, Timothy S. Coalson^e, Antonios Makropoulos^a, Jelena Bozek^f, Robert Wright^h, Andreas Schuh^a, Matthew Websterⁱ, Jana Hutter^g, Anthony Price^g, Lucilio Cordero Grande^g, Emer Hughes^g, Nora Tusor^g, Philip V. Bayly^c, David C. Van Essen^e, Stephen M. Smithⁱ, A. David Edwards^g, Joseph Hajnal^g, Mark Jenkinsonⁱ, Ben Glocker^a, Daniel Rueckert^a,

^a*Biomedical Image Analysis Group, Department of Computing, Imperial College London, London, United Kingdom*

^b*Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, MO, USA*

^c*Department of Mechanical Engineering and Material Science, Washington University in St. Louis, St. Louis, MO, USA*

^d*St. Lukes Hospital, Saint Louis, USA.*

^e*Department of Neuroscience, Washington University School of Medicine, St Louis, MO, USA*

^f*Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia*

^g*Centre for the Developing Brain, Division of Imaging Sciences, King's College London, London, United Kingdom.*

^h*Department of Biomedical Engineering, Division of Imaging Sciences, Kings College London, London, United Kingdom*

ⁱ*Centre for Functional Magnetic Resonance Imaging of the Brain, John Radcliffe Hospital, Oxford University, United Kingdom*

Abstract

The accurate alignment of brains is fundamental to the statistical sensitivity and spatial localisation of group studies in brain imaging, and cortical surface-based alignment is generally accepted to be superior to volume-based

*Corresponding author

Email address: ecr05@doc.ic.ac.uk (Emma C. Robinson)

¹These authors contributed equally

approaches at aligning cortical areas. However, human subjects have considerable variation in cortical folding, and in the location of cortical areas relative to these folds, which makes aligning cortical areas based on folding alone a challenging problem. The Multimodal Surface Matching (MSM) tool is a flexible spherical registration approach that enables accurate registration of surfaces based on a variety of different features. Using MSM, we have previously shown that using areal features such as resting state networks and myelin maps to drive cross-subject surface alignment improves group task fMRI statistics and map sharpness. However, the initial implementation of MSM's regularisation function did not penalize all forms of surface distortion evenly. In some cases, this allowed peak distortions to exceed neurobiologically plausible limits unless the regularisation strength was increased, in which case this prevented the algorithm from fully maximizing surface alignment. Here, we propose a new regularisation penalty, derived from physically relevant equations of strain (deformation) energy, and demonstrate that its use leads to improved and more robust alignment of multi-modal imaging data. In addition, since spherical warps incorporate projection distortions that are unavoidable when mapping from a convoluted cortical surface to the sphere, we also propose constraints to enforce smooth deformation of cortical anatomies. We test the impact of this approach for longitudinal modeling of cortical development for neonates (born between 32 and 45 weeks) and demonstrate that the proposed method increases the biological interpretability of the distortion fields and improves the statistical significance of population-based analysis relative to other spherical methods.

Keywords: Surface-based cortical registration; longitudinal registration;

neonatal brain development; Discrete Optimisation; biomechanical priors

1. Introduction

The cerebral cortex is a highly convoluted structure, with complex patterns of folding that vary considerably across individuals. Accurate cross-subject volumetric registration, in the face of folding variability, is far from straightforward as relatively small deformations in three dimensions risk matching opposing banks of cortical folds, or aligning brain tissue with cerebrospinal fluid. For this reason surface registration methods have been proposed, which constrain alignment to the 2D cortical sheet (Durrleman et al., 2009; Fischl et al., 1999b; Gu et al., 2004; Lombaert et al., 2013; Lyu et al., 2015; Robinson et al., 2014; Tsui et al., 2013; Wright et al., 2015; Yeo et al., 2010; Van Essen, 2005). These are generally accepted to have better performance at aligning cortical areas Glasser et al. (2016b)

Often, surface registration techniques have focused on the alignment of cortical convolutions. Examples include, spectral embedding approaches, which learn fast and accurate mappings between low-dimensional representations of cortical shapes (Lombaert et al., 2013; Orasanu et al., 2016b; Wright et al., 2015); Large Deformation Diffeomorphic Metric Mapping (LDDMM) frameworks, which learn vector flows fields between cortical geometries (Durrleman et al., 2009), to allow smooth deformation of cortical shapes (Durrleman et al., 2013); and spherical projection methods (Fischl et al., 1999b; Lyu et al., 2015; Van Essen et al., 2012; Yeo et al., 2010), which simplify the problem of cortical registration by projecting the convoluted surface to a sphere. All methods demonstrate clear advantages in terms of: improving the speed and accuracy of alignment (Lombaert et al., 2013; Wright et al., 2015; Yeo et al., 2010); increasing the correspondence of important features

on the cortical surface, such as brain activations (Fischl et al., 2008; Lombaert et al., 2015; Yeo et al., 2010); and providing a platform through which cortical shapes can be statistically compared (Durrleman et al., 2013; Orasanu et al., 2016b)

Ultimately, however, shape based alignment of the brain is limited as cortical folding patterns vary considerably across individuals, and correspond poorly with the placement of cortical architecture, function, connectivity, and topographic maps across most parts of the cortex (Amunts et al., 2000, 2007; Glasser et al., 2016a). For this reason several papers have been proposed to drive cortical alignment using 'areal' features (descriptors that correlate with the regional organisation of the human brain). These include Conroy et al. (2013); Frost and Goebel (2013); Nenning et al. (2017); Sabuncu et al. (2010), which drive spherical registration using features derived from functional Magnetic Resonance Imaging (fMRI), and Tardif et al. (2015), who register level-set representations of cortical volumes using combinations of geometric features and cortical myelin. Further, in Lombaert et al. (2015) and Orasanu et al. (2016a) spectral shape embedding approaches are extended to utilise broader feature sets, with Lombaert et al. (2015) adapting spectral alignment of the visual cortex to improve transfer of retinotopic maps, and Orasanu et al. (2016a) improving correspondence matching between neonatal feature sets by performing multimodal spectral embeddings of cortical shape and diffusion MRI (dMRI).

Most of these methods are optimised for alignment of specialised feature sets, whether that be fMRI time series (Sabuncu et al., 2010), functional connectivity patterns (Conroy et al., 2013; Nenning et al., 2017) or fixed

combinations of cortical geometry with fMRI/dMRI (Lombaert et al., 2015; Orasanu et al., 2016a). Furthermore, methods which enforce geometric constraints on the mapping (Durrleman et al., 2013; Lombaert et al., 2013, 2015; Orasanu et al., 2016b; Wright et al., 2015) are not well placed to address the known disconnect between the functional and morphological organisation of the human brain (Amunts et al., 2000; Nenning et al., 2017)

Therefore, in Robinson et al. (2014) we proposed Multimodal Surface Matching (MSM), a spherical deformation approach that enables flexible alignment of any type or combination of features that can be mapped to the cortical surface. MSM is inspired by DROP (Glocker et al., 2008) a discrete optimisation approach for non-rigid registration of 3D volumes. Like DROP, MSM has advantages in terms of reduced sensitivity to local minima (Glocker et al., 2011), and a modular optimisation framework. This means that any combination of similarity and regularisation terms can be used, which allows the framework to adapt to alignment of any type, or combination, of features provided that an appropriate similarity cost can be found. Accordingly, MSM has been used to drive alignment of a wide variety of different feature sets, including correlates of cortical folding, MR-based estimates of cortical myelination, resting-state network maps, and multi-modal combinations of folding and myelin (Božek et al., 2016; Glasser et al., 2016a,b; Harrison et al., 2015; Robinson et al., 2014). Experiments show that this also leads to improvements in correspondence of unrelated areal features including retinotopic maps (Abdollahi et al., 2014) and task fMRI (Glasser et al., 2016a; Robinson et al., 2014; Tavor et al., 2016).

One limitation of the original MSM implementation was that the choice

of discrete optimisation allowed only for first-order (pairwise) smoothness constraints. These was implemented by penalising the method from estimating very different transformations for neighbouring points on the surface. Unfortunately, this is suboptimal, as the penalty for doubling the size of a region is equal to that of reducing the area to zero. Further, it is not possible to optimally weight isotropic and anisotropic distortions. We therefore required a method for considering higher-order deformations, through comparing displacements of triplets of points; for this we implemented a method for higher-order discrete optimisation through clique reduction (Ishikawa, 2009, 2014), which has already been successfully applied to 2D registration in (Glocker et al., 2010)

A prototype of this framework proved fundamental to the methods used during developed of the HCP’s multi-modal parcellation (Glasser et al., 2016b,a). For (Glasser et al., 2016a), we proposed an triplet-based angular deviation penalty, that penalised change in the angles of each triangular mesh face. Unfortunately, this new regularization penalty measure also had drawbacks. In particular, it did not have any direct penalty for increasing or decreasing the size of a feature. With careful tuning of the regularization strength, based on elimination of neurobiologically implausible individual subject peak distortions, we were able to produce the publicly released MS-MAll aligned HCP results and the HCP’s multi-modal parcellation Glasser et al. (2016b,a). However, the need to eliminate neurobiologically implausible peak distortions limited the registration’s ability to maximize functional alignment. In particular, the method struggled to sufficiently penalize excessive distortion in regions where the individual topological layout of cortical

areas deviated from the group (Amunts et al., 2000; Gordon et al., 2017; Glasser et al., 2016a; Haxby et al., 2011; Wang et al., 2015). High regularization strength was required to prevent extreme distortions that would have led to issues with circularity and poor individual subject areal classifier performance (described in detail in Glasser et al. (2016a) Supplementary Methods sec 6.4).

In this paper we therefore present the high-order MSM framework with a new penalty inspired by the hyperelastic properties of brain tissue, to minimize surface deformations in a physically plausible way (Knutson et al., 2010); fully controlling both isotropic and anisotropic distortions, such that peak distortions can be robustly controlled in areas affected by topological variation or noise without requiring an excessively high regularization strength that decrease the overall alignment. We test this new strain-based regularization by re-optimizing the alignment of data from the adult Human Connectome Project (HCP), for both folding alignment (MSMSulc) and the multi-modal alignment protocol (MSMAll) described in Glasser et al. (2016a) (sec 6.4). This uses myelin maps, resting state network maps, and resting state visuotopic maps to align cortical areas.

A further limitation of the original MSM framework for longitudinal analyses is that use of spherical alignment complicates the neurobiological interpretation of deformation fields, since spherical projection distorts the relative separation of vertices between the 3D anatomical surface ('cortical anatomy') and the sphere, such that spherical regularisation has a varying influence on different parts of the cortical anatomy. These differences vary across the brain, and are size and shape dependent (Figure 1), implying spherical regu-

larisation has a varying influence on different parts of the cortical anatomy, and these effects vary across brains. In this paper, therefore, we propose an extension to the spherical alignment method, adapted from Knutsen et al. (2010), to deform points on the sphere but regularise displacements on the real anatomical surface meshes. This novel method retains the flexibility of the spherical framework (which allows registration of any type of feature that can be projected to the cortical surface), while harnessing spatial information from the anatomical surface to produce physically meaningful warps.

We test the approach on alignment of 10 longitudinally-acquired neonatal cortical surface data sets imaged twice between 31 and 43 weeks post-menstrual age and explore whether the deformation fields generated using the proposed method display improved correspondence (relative to other spherical alignment approaches) with expected growth trajectories over this developmental period.

The rest of paper is organised as follows: we first give an overview of the original spherical MSM method (sMSM), proposed in (Robinson et al., 2014) (sec. 2), and discuss the extension to higher-order regularisation penalty terms (sec. 3). Methods for approximating and regularising anatomical warps (aMSM) are then presented (sec. 4). Finally, experiments and results are presented on a feature sets derived from adult and developing Human Connectome Project (sec.6). Code and configuration files for running these experiments are available (https://www.doc.ic.ac.uk/~ecr05/MSM_HOCR_v2/); note, that the released HCP MSMA11 data used an older version of MSMA11 (see above, <https://db.humanconnectome.org>), however the MSMSulc and MSMA11 pipelines based on strain-based regulariza-




Figure 1. Metric distortions in the relative spacing of vertices occur as a result of the projection from anatomy to a sphere. These change across the surfaces and between brains. Differences are particularly obvious for longitudinally acquired data. Shown here: A) White matter surfaces extracted from the same subject at 34 weeks post-menstrual age (PMA, top row), and 44 weeks PMA (bottom row) are projected to a sphere; B) metric distortions estimated in terms of isotropic expansion of mesh faces ($\log_2 J$, for J defined in eq. 4), shown aligned and resampled to the 44 week subject (inflated brain view); C) metric distortion difference between time points

tion is publicly available (<https://github.com/Washington-University/Pipelines>) and will be used in follow up HCP projects on development, aging and human disease and preliminary dHCP data can be found from <https://data.developingconnectome.org> (version 1.1).

2. Multimodal Surface Matching

We begin with an overview of the original MSM method, first proposed in Robinson et al. (2014). In this framework, we seek alignment between two anatomical surfaces, each projected to a sphere, through the procedures outlined in Fischl et al. (1999a). Here, anatomical surfaces represent tessellated meshes fit to the outer boundary of a white matter tissue segmentation (i.e. the gray/white surface). These are expanded outwards to the outer grey

matter (or pial) boundary, and a midthickness surface is generated half way between white and pial boundaries. Separately, white matter surfaces are expanded to generate a smooth inflated cortical surface, from which points are then projected to a sphere. This is done in such a way as to minimise metric distortions i.e. minimise the overall change in area of triangular mesh faces during the transformation, Fischl et al. (1999a). Further, during inflation, indexed vertices (or points) in each surface space retain one-to-one correspondence, such that each index represents the same cortical location across white, midthickness, pial, inflated and spherical surfaces.

The goal of MSM registration is to identify spatial correspondences between two spheres so as to improve the overlap of the surface geometry and/or functional properties of the cortical sheet from which the spheres were derived. Spheres may represent corresponding hemispheres (left or right) from: two different subjects; one subject and a population average template; or the same subject imaged at two different ages. During registration, vertex points on one (source) sphere are moved until the surface properties on that sphere better agree with those of the second (target) sphere. Due to the vertex correspondence between sphere and surface anatomy, this also implicitly derives correspondences for the cortical anatomy.

Let SSS be the source spherical surface with initial coordinates \mathbf{x} , and TSS be the target spherical surface with coordinates \mathbf{X} , where $\mathbf{x}, \mathbf{X} \in \mathbb{S}^2$. Let SAS be the anatomical (white/midthickness/pial) surface representation of SSS with coordinate \mathbf{y} and TAS be the fixed anatomical surface representation of TSS with coordinates \mathbf{Y} , where $\mathbf{y}, \mathbf{Y} \in \mathbb{R}^3$. Let the moving source spherical surface be represented as MSS with coordinates \mathbf{x}' , and the

resulting deformed anatomical configuration as (*DAS*) with coordinates \mathbf{y}' ($\mathbf{x}' \in \mathbb{S}^2, \mathbf{y}' \in \mathbb{R}^3$). Anatomical surfaces may be white, midthickness, pial or inflated surfaces, and each surface is associated with multimodal feature sets \mathbf{M} (fixed) and \mathbf{m} (moving $\mathbf{M}, \mathbf{m} \in \mathbb{R}^N$). These represent any combination of N features describing cortical folding, brain function (such as resting state networks), cortical architecture, or structural connectivity. The set of surfaces involved are shown in Fig. 2.

MSM employs a multi-resolution approach. In this, a sequence of spherical, regularly-sampled, control-point grids $(\mathbf{G}_D)_{D \in \mathbb{N}}$ are used to constrain the deformation of *MSS* (Fig. 2b). These are formed from regular subdivisions of icospheric, triangulated meshes, where the granularity of the D th control-point grid increases at each resolution, allowing features of the data to be matched in a coarse-to-fine fashion. Typically, at each resolution, the features from *MSS* and *FSS* are downsampled onto regular data grids MSS_D and FSS_D to speed processing. Resampling the data in this way may also reduce any impact that the meshing structures of *MSS* and *FSS* might have on the deformation. Final upsampling of the control-point warp to *MSS* is performed using barycentric interpolation.

At each resolution level, registration proceeds as a series of discrete displacement choices. At each iteration, points $\mathbf{p} \in \mathbf{G}_D$, are given a finite choice of possible locations on the surface to which they can move. The end points of each displacement are determined from a set of L vertex points defined on a regular sampling grid (Fig. 2c, purple crosses, see also Robinson et al. (2014)). Displacements are then defined in terms of a set of L rotation matrices $S^{\mathbf{p}} = \{\mathbf{R}^1, \mathbf{R}^2 \dots \mathbf{R}^L\}$, specified separately for each control point \mathbf{p} . These

rotate \mathbf{p} to the sample vertex points by angles expressed relative to the centre of the sphere.

The optimal rotation for each control point $\mathbf{R}_{\mathbf{p}} \in S^{\mathbf{p}}$ is found using discrete optimisation (Robinson et al., 2014; Glocker et al., 2008). This balances a unary data similarity term $c(\mathbf{R}_{\mathbf{p}})$ with a pairwise penalty term $V(\mathbf{R}_{\mathbf{p}}, \mathbf{R}_{\mathbf{q}})$, which encourages a smooth warp. The search for the optimal rotations can be defined as cost function (C) over all unary and pairwise terms as:

$$\min C = \sum_{\mathbf{p} \in \mathbf{G}_D} c(\mathbf{R}_{\mathbf{p}}) + \lambda \sum_{\mathbf{p} \in \mathbf{G}_D} \sum_{\mathbf{q} \in N(\mathbf{p})} V(\mathbf{R}_{\mathbf{p}}, \mathbf{R}_{\mathbf{q}}) \quad (1)$$

Here $\mathbf{q} \in N(\mathbf{p})$ represents all control points that are neighbours of \mathbf{p} , and λ is a weighting term that balances the trade off between accuracy and smoothness of the warp.

One advantage of the discrete framework is that cost functions do not need to be differentiable, and thus there are no constraints on the choice of data similarity term $c(\mathbf{R}_{\mathbf{p}})$, for example: correlation, Normalised Mutual Information (NMI), Sum of Square Differences (SSD), and alpha-Mutual Information (α -MI, useful for multi-modal alignments) (Neemuchwala, 2005; Robinson et al., 2014) can all be applied.

However, a limitation of the original discrete optimisation framework (Robinson et al., 2014) has been that it is limited to first-order regularisation constraints. This meant that smoothness constraints are imposed through pairwise regularisation terms, implemented by penalising differences between

proposed rotation matrices for neighbouring control-points:

$$V(\mathbf{R}_{\mathbf{p}}, \mathbf{R}_{\mathbf{q}}) := \|\log((\mathbf{R}_{\mathbf{p}}\mathbf{R}'_{\mathbf{p}})^T(\mathbf{R}_{\mathbf{q}}\mathbf{R}'_{\mathbf{q}}))\|_F^2 \quad (2)$$

Here $\|\cdot\|_F$ represents the Frobenius norm, and $\mathbf{R}'_{\mathbf{p}}$ represents the full rotation of the control-point \mathbf{p} , accumulated over previous labelling iterations. In what follows we refer to this regularisation framework as *sMSM_{PAIR}*.

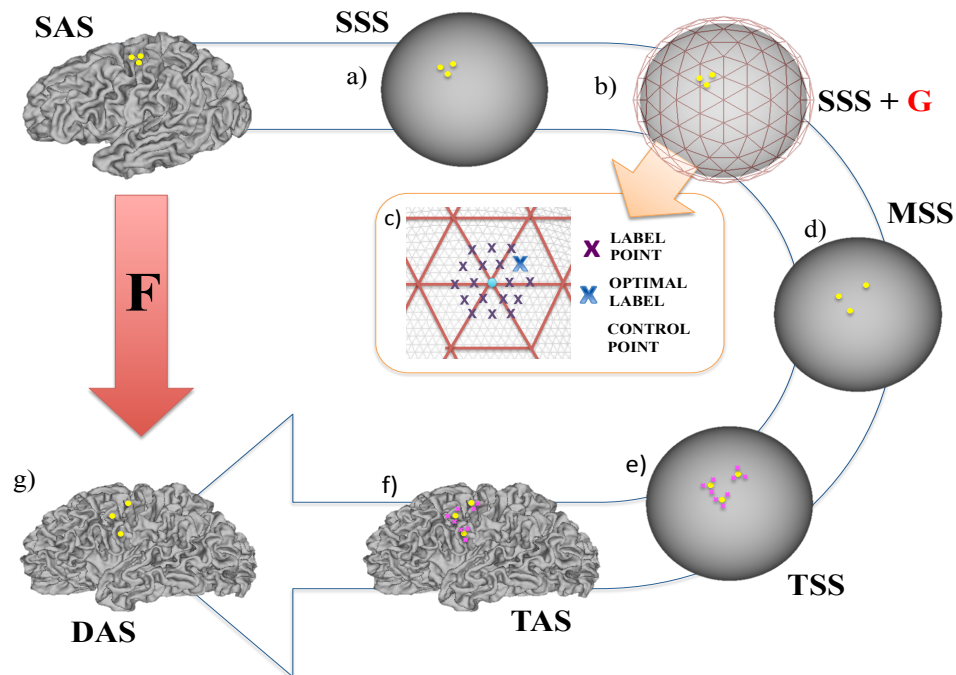


Figure 2. Projecting cortical anatomy through spherical warps. The figure follows the displacement of three (yellow) points on the source white anatomical surface (SAS), via the moving source spherical surface (MSS), into a new configuration on the target white anatomical surface (TAS), where source and target represent the left hemispheres of two different subjects aged 38 ± 1 week PMA. Steps: a) Vertex correspondence between the source sphere (SSS) and anatomy (SAS) means that points form triplets on both surfaces; b) Control point grids (G , red) constrain the deformation of SSS within a discrete optimisation scheme (orange box). c) Each control-point (blue dot) can move to a finite number of possible positions on the surface (purple crosses). The optimal displacement (blue cross) improves feature map similarity whilst constraining deformations to be smooth; d) The displaced spherical surface configuration MSS is estimated from G using barycentric interpolation (Eq. 7); e) Barycentric correspondences are learnt between vertices on MSS (yellow dots) and TSS (pink crosses; Eq. 9); f) Weights (calculated during step e) are applied to the equivalent points on the target anatomical surface TAS (Eq. 10); creating g) a deformed anatomical surface configuration (DAS), which has the mesh topology of the source SAS surface, but the shape of the target anatomical surface (TAS). Through this a transformation \mathbf{F} can be estimated between SAS and DAS .

3. Higher-order Smoothness Constraints and Strain-based Regularization

As described in the introduction, it is helpful to compute both similarity and regularization using whole triangles, rather than pairs of vertices, and to fully model isotropic and anisotropic surface distortions. In this paper, we therefore propose to take advantage of recent advances in discrete optimisation (Ishikawa, 2009, 2014) that allow adoption of higher-order smoothness constraints.

This necessitates generalisation of the original cost function (equation 1) to allow for terms (formally known as cliques) to vary in size:

$$\min C = \sum_{c_1 \in C_D} c(R_{c_1}) + \lambda \sum_{c_2 \in C_R} V(R_{c_2}) \quad (3)$$

Here C_D represents cliques used for estimation of the data similarity term, C_R represents regularisation cliques, and R_{c_1} , R_{c_2} represent the sub-set of rotations estimated for each clique. This represents a highly modular framework where any combination of similarity metric and smoothness penalty can be used, provided they can be discretised as a sum over clusters of nodes the graph. In this paper we focus on two new triplet terms:

- **Triplet Regularisation** $V(R_{c_2})$: We propose a new regularisation term derived from biomechanical models of tissue deformation. This term is inspired by the strain energy minimization approach used in Knutsen et al. (2010), which constrains the strain energy density (W_{pqr}) of locally affine warps \mathbf{F}_{pqr} , defined between the control-point mesh faces $\mathcal{T} = \{\mathbf{p}, \mathbf{q}, \mathbf{r}\}$.

Specifically, \mathbf{F}_{pqr} represents the 2D transformation matrix, or deformation gradient, for $\mathcal{T} = \{\mathbf{p}, \mathbf{q}, \mathbf{r}\}$ after projection into the tangent plane, fully describing deformations on the surface. The eigenvalues of \mathbf{F} represent principal in-plane stretches, λ_1 and λ_2 , such that relative change in area may be described by $J = \lambda_1\lambda_2$ and relative change in shape (aspect ratio) may be described by $R = \lambda_1/\lambda_2$. Areal distortion is traditionally defined as $\log_2(Area_1/Area_2) = \log_2 J$. Here we introduce an equivalent term for shape distortion, defined as $\log_2 R$. To penalize against both types of deformations, we define strain energy density using the form:

$$W_{pqr} = \frac{\mu}{2}(R^k + R^{-k} - 2) + \frac{\kappa}{2}(J^k + J^{-k} - 2) \quad (4)$$

where k is defined as any integer greater than or equal to 1. This proposed form meets all criteria of a hyperelastic material (a class often used to characterize biological soft tissues including brain). As described for other hyperelastic materials, shear modulus, μ , penalizes changes in shape and bulk modulus, κ , penalizes changes in volume, or in the case of 2D, area. Our improved form ensures that expansion ($J > 1$) and shrinkage ($J < 1$) are penalized equally in log space (critical to avoid imbalances in surface registration such as uncontrolled shrinkage -see introduction) and that changes in shape (R) or area (J) are penalized by the same function. Additionally, this regularization function allows the optimization of the trade off between areal distortion and shape distortion. The strain energy penalty is implemented as:

$$V_{STR}(\mathbf{R}_p, \mathbf{R}_q, \mathbf{R}_r) := W_{pqr}^2 \quad (5)$$

Note, for the case of $k = 1$, it can be shown that this form is equivalent to a modified, compressible Neo hookean material, similar to the original form used in Knutsen et al. (2010). See Supplementary Material for more details and formal justification of the proposed strain energy form.

- **Triplet Likelihood $V(R_{c_2})$:** First proposed for registration in Glocker et al. (2010), triplet likelihoods were introduced as a means of setting up a well-posed image matching problem, since (for spherical registration) a two-dimensional displacement must be recovered from a one-dimensional similarity function. As in Glocker et al. (2010) we implement triplet likelihood terms as correlations (CC) between patches of data: defined as all data points which overlap with each control-point mesh face triplet:

$$\psi(\mathbf{R}_p, \mathbf{R}_q, \mathbf{R}_r) = CC(\mathbf{m}_{pqr}, \mathbf{M}_{pqr}) = 1 - \frac{cov((\mathbf{m}_{pqr}, \mathbf{M}_{pqr}))}{\sigma_{\mathbf{m}_{pqr}} \sigma_{\mathbf{M}_{pqr}}} \quad (6)$$

where, \mathbf{m}_{pqr} is the sub-matrix of features from \mathbf{m} , which correspond to points from the moving mesh MSS that move with the control point triplet $\mathcal{T} = \{\mathbf{p}, \mathbf{q}, \mathbf{r}\}$. \mathbf{M}_{pqr} represent the overlapping patch in the fixed mesh space TSS . Features \mathbf{M}_{pqr} are resampled onto MSS using adaptive barycentric resampling (Glasser et al., 2013), and $\sigma_{\mathbf{m}_{pqr}}$ and $\sigma_{\mathbf{M}_{pqr}}$ represent the variances of each patch.

4. Anatomical Regularisation (aMSM)

4.1. Inferring Anatomical Correspondences

Using the known vertex correspondence between the fixed sphere (*TSS*) and its anatomical representation (*TAS*, Fig. 2e) a deformed anatomical surface configuration (*DAS*) can be found for the moving surface in the following steps:

1. The coordinates (\mathbf{x}') of the moving source sphere (*MSS*) are found by interpolating coordinates from the control point grid triplet $\mathcal{T}_{\mathbf{x}} = \{\mathbf{p}, \mathbf{q}, \mathbf{r}\}$ that overlaps with \mathbf{x} , such that :

$$\mathbf{x}' = \eta(\mathbf{x}, \mathbf{p}; \mathbf{q}, \mathbf{r})\mathbf{R}_{\mathbf{p}}\mathbf{R}'_{\mathbf{p}}\mathbf{p} + \eta(\mathbf{x}, \mathbf{q}; \mathbf{r}, \mathbf{p})\mathbf{R}_{\mathbf{q}}\mathbf{R}'_{\mathbf{q}}\mathbf{q} + \eta(\mathbf{x}, \mathbf{r}; \mathbf{p}, \mathbf{q})\mathbf{R}_{\mathbf{r}}\mathbf{R}'_{\mathbf{r}}\mathbf{r} \quad (7)$$

Here, $\mathbf{R}_{\mathbf{p}}\mathbf{R}'_{\mathbf{p}}$ represents the combined rotation of the control point over all iterations, and $\eta(\cdot)$ is a barycentric interpolation function:

$$\eta(\mathbf{x}, \mathbf{p}; \mathbf{q}, \mathbf{r}) = \frac{[\mathbf{x}\mathbf{q}\mathbf{r}]}{[\mathbf{p}\mathbf{q}\mathbf{r}]} \quad (8)$$

where [...] represents triangle area.

2. Barycentric correspondences are found between the moving spherical surface configuration (*MSS*) and the fixed sphere (*TSS*). These are used to define a set of vertex indices, and corresponding weights, sufficient for resampling coordinates from the fixed topology onto the moving topology

$$\eta_{M \rightarrow F}(\mathbf{x}', \mathbf{X}_i; \mathbf{X}_j, \mathbf{X}_k) = \frac{[\mathbf{x}'\mathbf{X}_j\mathbf{X}_k]}{[\mathbf{X}_i\mathbf{X}_j\mathbf{X}_k]} \quad (9)$$

For $\mathcal{T}_{\mathbf{x}'} = \{\mathbf{X}_i, \mathbf{X}_j, \mathbf{X}_k\}$, where $\mathcal{T}_{\mathbf{x}'}$ is a triplet of points on the fixed surface (*FSS*) that overlaps \mathbf{x}' .

3. The indices and weights found in Eq. 9 are used to project the moving surface mesh topology onto the fixed anatomical surface. We call this the deformed anatomical surface (*DAS*) as this implements the warp that is implied through the allocation of point-wise correspondences during the spherical warp:

$$\mathbf{y}' = \eta_{M \rightarrow F}(\mathbf{x}', \mathbf{X}_i; \mathbf{X}_j, \mathbf{X}_k) \mathbf{Y}_i + \eta_{M \rightarrow F}(\mathbf{x}', \mathbf{X}_j; \mathbf{X}_k, \mathbf{X}_i) \mathbf{Y}_j + \eta_{M \rightarrow F}(\mathbf{x}', \mathbf{X}_k; \mathbf{X}_i, \mathbf{X}_j) \mathbf{Y}_k \quad (10)$$

4.2. Implementing aMSM Regularisation

Using the higher-order constraints formulated in section 3, the evolution of (*DAS*) can be controlled by replacing spherical control-point triplets, with anatomical triplets in Equation 4. In the simplest case this can be achieved by introducing low-resolution anatomical surfaces at the resolution of the control-point grid G_D . In this way a low-resolution deformed anatomical configuration DAS_D can be determined from correspondences found between the control-point grid and a fixed low-resolution target sphere, TSS_D , and associated low-resolution anatomical grid TAS_D , with coordinates $\mathbf{y}, \mathbf{Y} \in \mathbb{R}^3$ and $\mathbf{X}^D \in \mathbb{S}^2, \mathbf{Y} \in \mathbb{R}^3$. Coordinates $\mathbf{y}' \in \mathbb{R}^3$ are then estimated using:

$$\mathbf{y}'^D = \sum_{i=0}^2 \eta_{M \rightarrow T}(\mathbf{p}, \mathbf{X}_i^D \mid \mathcal{T}_{\mathbf{p}} \setminus \{\mathbf{X}_i^D\}), \mathbf{Y}_i^D \quad (11)$$

This is shorthand for the notation in Eq 10, where $\mathcal{T}_{\mathbf{p}} \setminus \{\mathbf{X}_i^D\}$ represents the points in the triplet $\mathcal{T}_{\mathbf{p}}$ excluding \mathbf{X}_i^D . Therefore, transformations (\mathbf{F} in Fig. 2) are assessed by comparing matching triplets between \mathbf{y}^D and \mathbf{y}'^D . Nevertheless, it is important to note that the discrete displacement space continues to be estimated on the sphere. This allows registration to achieve

lower anatomical distortions without sacrificing the quality of the alignment. We refer to this adaptation as anatomical MSM (*aMSM*).

5. Implementation details

The proposed registration cost functions are non-convex, due to local minima in the similarity function and nonlinear penalty terms. This means that they cannot be solved by conventional discrete solvers such as α -expansion (which is used in graph cuts (Boykov et al., 2001)) as this requires that pairwise terms be submodular (meet the triangle inequality). Instead, it is necessary to use methods that account for non-submodularity such as FastPD (Komodakis and Tziritas, 2007; Komodakis et al., 2008) and QPBO (Rother et al., 2007). However, these methods only solve pairwise MRF functions.

In order to account for triplet terms we adopt the approach of Ishikawa (2009, 2014). This allows reduction of higher order terms either by: A) addition of auxiliary variables (for example by reducing a triplet to three pairwise terms (Ishikawa, 2009); or B) by reconfiguration of the polynomial form of the MRF energy, until the high-order function can be replaced by a single quadratic (Ishikawa, 2014). We present results using the latter version, known as Excludable Local Configuration (ELC). This has the advantage that, provided an ELC can be found, there is no increase in the number of pairwise terms, which has some impact on the computational time.

Once terms are reduced, optimisation proceeds as solutions to a series of binary label problems, where results for the full label space are obtained using the hierarchical implementation of the fusion moves technique (Lempitsky et al., 2010), as described in Glocker et al. (2010). In each instance, the

reduction is passed to the FastPD solver (Komodakis and Tziritas, 2007; Komodakis et al., 2008) for optimisation.

For effective *aMSM* implementation, choices have to be made with regards to how far the anatomy should be reasonably downsampled. Resampling of the anatomical surface is performed by barycentric interpolation, using correspondences between the initial source sphere (*SSS*) and regular, low-resolution icospheric spherical grids (Supplementary Material). Where resolution of *SAS^D* exceeds that of the control-point grid, regularisation for each control-point triplet is calculated by averaging distortions for every higher-resolution triplet in *SAS^D* that falls within the area of that control-point grid face. This results in a trade-off between run time and accurate capture of the true anatomical distortion (see Supplementary Material). For every increase in anatomical mesh resolution relative to the control point grid, strain calculations better represent that of the native deformation, but number of strain calculations increases by a factor of four.

In general, estimation of triplet energies and reduction through ELC and binary FastPD slows the run time relative to the original MSM approach. To reduce some of the impact the control point grids are no longer reset after each iteration. Instead the source mesh and control grid incrementally deform together and all neighbourhood relationships are learned once at the beginning of each resolution level. To prevent folding of the mesh during alignment a weighting penalty is placed on the regularisation cost that severely penalises flipping of the triangular faces. This is done to ensure the final transformation is smooth and invertible.

Finally, to improve convergence and allow for smoother warps, modi-

fications are also made to the rescaling of the discrete label space at each iteration. In the original framework the label space switches between the vertices and barycentres of a regular sampling grid (Robinson et al., 2014). In the new framework the discrete displacement vectors \mathbf{d}^l_p are instead rescaled by $0.8\times$ their original length over 5 iterations, before being reset to their original lengths.

6. Experimental Methods and Results

We test the framework on real data collected as part of the adult Human Connectome Project (HCP²) and Developing Human Connectome Project (dHCP³) to assess the impact of the proposed strain regulariser on both spherical and anatomical deformations. In all experiments results are compared for the higher-order MSM registration framework (MSM_{STR}) and the original form ($sMSM_{PAIR}$). Where feasible MSM was also compared against FreeSurfer (FS : arguably the most commonly used tool for cortical surface alignment) and Spherical Demons (SD), which is diffeomorphic on the sphere. Note, standard implementations of FreeSurfer and Spherical Demons are not capable of multimodal or multivariate alignment.

6.1. Cohorts

A sub-set of 28 subjects for the full HCP cohort (1200 subjects) were selected for parameter optimisation of the multimodal alignment protocol laid out in Glasser et al. (2016a). Here registration was driven using a com-

²<http://www.humanconnectome.org/>

³<http://www.developingconnectome.org/>

combination of features reflecting myelin maps Glasser and Van Essen (2011), 34 well-defined cortical surface resting-state spatial maps, and 8 visuotopic features (reflecting topographic organisation of functional connectivity in the visual cortex; see Glasser et al. (2016a) Supplementary Methods sec 2). The sub-set of subjects was selected to reflect a spectrum of sources of variance in the HCP data, including data sets with unusual functional topology (see Glasser et al. (2016a) Supplementary Results sec 1.3); or low signal-to-noise ratio (SNR).

In the second experiment a subset of dHCP subjects scanned (twice) longitudinally were used to explore longitudinal alignment of developing cortical shapes. This sub-set was selected specifically to contain all subjects scanned twice within 32.66 ± 1.22 weeks PMA (first scan) and 41.47 ± 1.61 weeks PMA (second scan) in order to allow straight-forward comparison of the deformations across subjects.

6.2. Data

Acquisition of HCP data was performed on a Siemens 3T Skyra platform, using a 32-channel head coil and MPRAGE (T1w) and SPACE (T2w) sequences (Glasser et al., 2013). Isotropic structural image acquisitions were acquired at $0.7mm^3$. Functional imaging data was acquired with multiband (factor 8) 2mm Gradient-Echo EPI sequences (Moeller et al., 2010). Four resting state fMRI (rfMRI) scans were acquired (two successive 15 minutes scans in each of two sessions) (Smith et al., 2013). Seven task-fMRI (tfMRI) experiments were also conducted, including: working memory, gambling, motor, language, social cognition, relational, and emotional tasks; tfMRI scans were acquired after the rfMRI scans in each of two hour-long sessions on

separate days (Barch et al., 2013). Additional details regarding specific acquisition parameters and task protocols are available in Barch et al. (2013) and Smith et al. (2013), as well as on the HCP website⁴. Generation of surface meshes and associated shape features were carried out using HCP Structural Pipelines (Glasser et al., 2013).

dHCP data was acquired at St. Thomas Hospital, London, on a Philips 3T scanner using a 32 channel dedicated neonatal head coil (Hughes et al., 2016). To reduce the effects of motion, T2 images were obtained using a Turbo Spin Echo (TSE) sequence, acquired in two stacks of 2D slices (in sagittal and axial planes), using parameters: TR=12s, TE=156ms, SENSE factor 2.11 (axial) and 2.58 (sagittal). Overlapping slices (resolution (mm) $0.8 \times 0.8 \times 1.6$) were acquired to give final image resolution voxels $0.8 \times 0.8 \times 0.8\text{mm}^3$ after motion corrected reconstruction, combining Cordero-Grande et al. (2016); Kuklisova-Murgasova et al. (2012). T1 images were acquired using an IR-TSE (Inversion Recovery Turbo Spin Echo) sequence at the same resolutions with TR=4.8s, TE=8.7ms, SENSE factor 2.26 (axial) and 2.66 (sagittal). All images were reviewed by an expert paediatric neuroradiologist and checked for possible abnormalities. Generation of surface meshes and associated shape features were carried out using dHCP Structural Pipelines (Makropoulos et al., 2017).

6.3. Strain Parameter Optimisation

Higher-order MSM (MSM_{STR}) was run using: tri-clique data terms, bulk modulus (κ) of 1.6, shear modulus (μ) of 0.4 (i.e. a 4 to 1 ratio), and

⁴<http://www.humanconnectome.org/documentation/Q1/imaging-protocols.html>

$k = 2$. These parameters were optimised for multimodal alignment of adult HCP data sec. 6.4, where bulk and shear moduli were optimized for an exponent of $k = 2$, keeping $\kappa + \mu = 2$ and regularisation λ constant), until alignment was optimised. These values were kept throughout the paper. Experiments on the influence of these parameters for longitudinal alignment of cortical anatomies (run using $aMSM_{STR}$, Supplementary Material sec. 2) show that, in general, results are robust over a range of parameters. Further, the parameters chosen for adult multimodal data are also optimal for both adult human folding alignment and longitudinal folding alignment. This suggests the strain parameters generalise well.

6.4. Multimodal alignment of adult HCP data

In the first experiment, $sMSM_{STR}$ was compared against $sMSM_{PAIR}$, SD , and FS for alignment of task fMRI data from a sub-set of 28 subjects from the HCP project. In this instance, anatomical regularisation was not used lest it prejudiced the solution towards alignment of cortical folding patterns, as these may not consistently reflect areal features across large parts of the brain across subjects. MSM methods were compared against FS and SD , run using their default settings (cortical folding alignment only) because FS registration is fixed and immutable, and the current implementation of SD allows only for alignment of univariate features.

MSM was run in two stages: first registration was initialised using constrained alignment of cortical folds, as described in Robinson et al. (2014) as the MSMSulc protocol, optimized in order to maximize task fMRI alignment. Then alignment of areal features was refined using what has become known as MSMAII Glasser et al. (2016a). In this multimodal registration

features containing myelin maps Glasser and Van Essen (2011), resting-state networks (RSNs) and visuotopic features were used to drive alignment to a group average template.

MSMSulc was run using strain-based regularisation ($sMSM_{STR}$) over three control-point (CP) grid resolutions with CP resolution: $CP_{res}=162, 642, 2542$; and features sampled to regularly-spaced grids ($MSSD, FSSD$) at resolutions: $DP_{res}=2542, 10242, 40962$. Regularisation strength was controlled through $\lambda = 10, 7.5, 7.5$. Features were variance normalised but no smoothing of the data was performed. Notably, the re-optimized strain-based MSMSulc substantially outperforms both the pairwise MSMSulc and the FreeSurfer registration (Table 1).

MSMAll was optimised three times, once using strain-based higher-order regularisation and likelihood terms ($sMSM_{STR}$) and twice using pair-wise regularisation $sMSM_{PAIR}$. In each case, common parameters between the methods were fixed; registration was run over three control-point grid resolutions: $CP_{res}=162, 642, 2542, DP_{res}=2542, 10242, 40962$; features were variance normalised and there was no smoothing of the data. Regularisation of $sMSM_{STR}$ was optimised in order to maximise alignment of tfMRI (appraised through estimates of cluster mass - defined below), subject to edge distortion mean across subjects and surfaces not exceeding that for FreeSurfer. Then $sMSM_{PAIR}$ was optimised twice: in the first case to achieve comparable peak edge distortions to $sMSM_{STR}$ (this keeps peak distortions well controlled, but limits the amount of registration that can occur, see introduction) , and in the second case to match comparable mean values of edge distortions (this allows a similar amount of registration to occur, but

peak distortions may exceed neurobiologically plausible thresholds).

Following registration tfMRI timeseries on the subjects' native meshes were resampled according to the registration to the standard mesh. Improvements in alignment were assessed via comparisons of the group mean activation maps (obtained using mixed effects FLAME Woolrich et al. (2009)) both qualitatively and quantitatively, via cluster mass, calculated using the following formula: $CM = \sum_{i \in S} |z(\mathbf{x}_i)| A(\mathbf{x}_i)$. Here x_i is a vertex coordinate, $z(\mathbf{x}_i)$ is the statistical value at this coordinate, $A(\mathbf{x}_i)$ is the area associated with this vertex (calculated from a share of the area of each mesh triangle connected to it in the mid-thickness surface), and S is the set of vertices where $|z(\mathbf{x}_i)| > 5$, and the threshold of 5 was chosen to be approximately equivalent to a two-tailed Bonferroni correction. The cluster mass measure reflects both the size of the super-threshold clusters and the magnitude of the statistical values within them.

Distortions are reported in terms of absolute values for: areal distortions ($\log_2 J$, Eq. 4), shape distortions ($\log_2 R$), and edge distortions. Edge distortions are estimated from the relative change in length of edges between neighbouring vertices in the mesh: $\log_2 \frac{L_2}{L_1}$, where L_2 is edge length following registration, and L_1 is length before, and are reported per vertex by taking the average values for all connected edges. These maps reflect a univariate summary of changes to area and shape and thus were used during optimisation

Results in Table 1 and Fig. 3 demonstrate that, as expected, multimodal alignment of tfMRI data (MSMAll) significantly improves the sharpness and peak values of the group z-statistics, relative to cortical folding based align-

ment (SD , FS , $MSMSulc$). This results in a 20.97% increase in total cluster-mass for MSM_{All} run with $sMSM_{STR}$ relative to $MSMSulc$, a comparable increase of 22.42% over SD , and a 27.31% increase over FS . Fig. 6 displays the spread of improvements across all contrasts within each task as improvement relative to FS , again $sMSM_{STR}$ outperforms all other methods.

In terms of folding-based methods, both $sMSM_{STR}$ and SD have much lower and smoother distortions than FS ; although it is important to note both are optimised for alignment of areal features Robinson et al. (2014); Yeo et al. (2010), whereas FS is optimised for alignment of cortical folds. $MSMSulc$, achieves marginal improvements over SD , with slight increases in $tMRI$ cluster-mass obtained from deformations with lower isotropic distortions, and similar edge distortions. Fig. 5 shows $sMSM_{STR}$ edge distortions are dispersed across the whole of the surface whereas SD alignment results in peaks of high distortions.

For multimodal alignment, optimising the original form of MSM to achieve comparable mean distortions does achieve comparable improvements in cluster mass to $sMSM_{STR}$ but leads to extensive patches of extreme distortions across the cortical averages (Fig. 5, Fig 5 D), with peak values for $\log_2 J$ and $\log_2 R$ exceeding that of $sMSM_{STR}$ by 128% and 63% respectively. These levels of distortions are highly non plausible given known ranges of regional variation from previous studies (Van Essen, 2005).

When distortions are instead matched for peak strains, improvements in total cluster mass observed for $sMSM_{STR}$ are 3.86% above the gains obtained with the original form of MSM (Table 1) showing that MSM_{All} does not reach its full potential for aligning the data. This results in sharper $tMRI$

group task maps (Fig. 3). Indeed, out of an abundance of conservatism we do not exceed the field standard edge distortion of FreeSurfer, however the strain-based MSM is capable of being pushed harder without peak distortions getting out of control. Unlike with MSMAll pairwise, with the strain-based MSMAll one can choose the desired level of distortion without having issues with runaway peak distortions, making the algorithm both more robust and easier for less neuroanatomically experienced users to use.

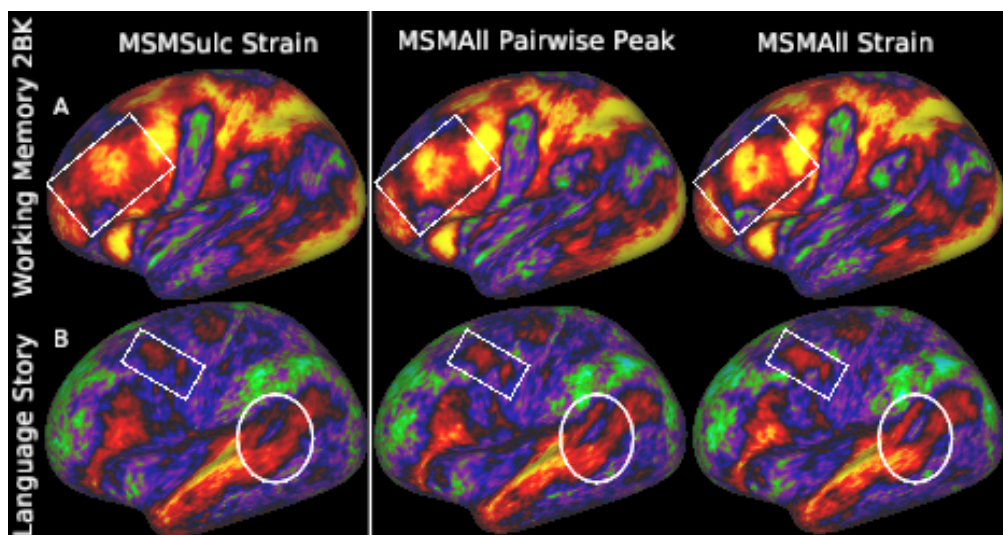


Figure 3. Comparison of group Z-statistic spatial maps following folding alignment (MSM-Sulc, run with $sMSM_{STR}$) and alignment driven by multimodal features (MSMAll, run with $sMSM_{STR}$ and $sMSM_{PAIR}$, matched for peak strains) for: a) a working-memory contrast (2BK) and b) a language task (Story). White boxes highlight improvements in sharpness of the contrast in the areas of the Dorsal Lateral Pre-Frontal cortex (A); region 55b (B) and in the temporal lobe (B)

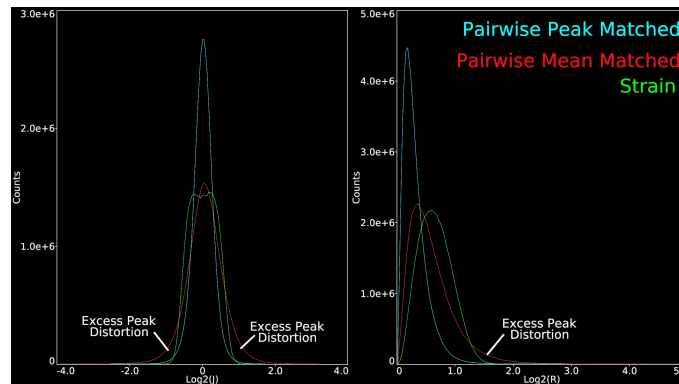


Figure 4. Histogram plots comparing MSMall distortions. $sMSM_{PAIR}$ generates long tailed distributions with excessive peak distortions

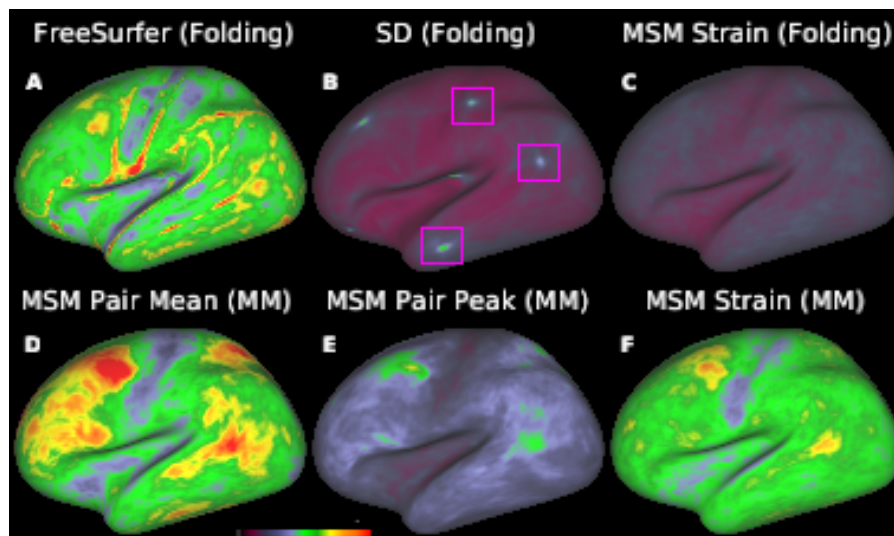


Figure 5. Mean edge distortion maps, averaged across all surfaces. Top row) distortions for folding based alignment only, peak boxes highlight hot spots of edge distortions for SD method; Bottom row) multi-modal (MM) alignments: MSM Pair Mean (MSMall run with $sMSM_{PAIR}$ optimised to achieve comparable mean strains to $sMSM_{STR}$); MSM Pair Peak (MSMall run with $sMSM_{PAIR}$ optimised to achieve comparable peak strains to $sMSM_{STR}$); MSM Pair Strain (MSMall run with $sMSM_{STR}$)




Figure 6. Bar chart of mean cluster mass statistics across HCP task categories for different methods. Note, only pure contrasts are included.

| | Cluster Mass | Edge distortion | | Areal distortion | | Shape distortion | |
|---|----------------|-----------------|-------|------------------|-------|------------------|-------|
| | | Mean | Max | Mean | Max | Mean | Max |
| <i>FS</i> | 3212876 | 0.271 | 2.564 | 0.377 | 4.879 | 0.698 | 6.752 |
| <i>SD</i> | 3341178 | 0.074 | 0.520 | 0.124 | 1.087 | 0.147 | 0.884 |
| MSMSulc (<i>sMSM_{PAIR}</i>) | 3019301 | 0.089 | 0.486 | 0.142 | 0.879 | 0.183 | 1.040 |
| MSMSulc (<i>sMSM_{STR}</i>) | 3381148 | 0.089 | 0.268 | 0.102 | 0.525 | 0.235 | 1.151 |
| MSMAll (<i>sMSM_{PAIR}</i> , peak matched) | 3938268 | 0.139 | 0.972 | 0.209 | 2.561 | 0.311 | 2.824 |
| MSMAll (<i>sMSM_{PAIR}</i> , mean matched) | 4022388 | 0.271 | 1.918 | 0.387 | 3.939 | 0.604 | 4.901 |
| MSMAll (<i>sMSM_{STR}</i>) | 4090190 | 0.261 | 0.919 | 0.306 | 1.731 | 0.662 | 3.015 |

Table 1. Peak distortions and cluster mass estimates following alignment of HCP tfMRI data using each of the proposed methods.

6.5. Longitudinal Registration

In the second experiment we explore the impact of anatomical warp regularisation for within-subject longitudinal alignment of 10 different neonatal subjects, each scanned twice within two specific time points (TPs): 32.66 ± 1.22 weeks PMA (TP1) and 41.47 ± 1.61 weeks PMA (TP2). This selection criterion was made by clustering all longitudinally scanned subjects (37 at time of writing) into groups with similar TPs, such that the resulting defor-

mations may be directly compared. This section presents results from the group with the biggest scan separation. Over this time period cortical geometry and the pattern of distortions resulting from the spherical projection changes dramatically (Fig. 7 a). Since the relationship between structural and functional organisation of the brain remains consistent over time, it is sufficient to align cortical geometry within a longitudinal experiment. Accordingly, in this experiment, registration was driven using mean curvature (Fig. 7 c).

Results are compared for: the original MSM spherical framework ($sMSM_{PAIR}$); spherical MSM with higher-order strain regularisation ($sMSM_{STR}$); anatomical MSM with strain regularisation ($aMSM_{STR}$); and Spherical Demons (SD). All registrations were run using the TP1 surface as the target of the registration, and TP2 as the moving surface as this was found to generate the most accurate alignment, where this was judged in terms of correlation between features sets, relative to total surface distortion. This is likely because it is a better-defined problem to register a more complex surface to a simpler surface than the reverse. Deformations in the direction TP1 \rightarrow TP2 were then obtained by inverting the transformation.

SD was run using its default parameterisation and represents a baseline for smooth diffeomorphic spherical alignment. All MSM registrations were optimised over 4 resolution levels with mesh resolutions: $CP_{res}=162, 642, 2542, 10242$, $DP_{res}=10242, 10242, 40962, 40962$; variance normalisation and smoothing was applied $\sigma_{in} = \sigma_{ref} = 6, 4, 2, 1$. $aMSM_{STR}$ was parameterised to penalise distortions of the midthickness surfaces, using anatomical mesh resolutions (DAS_D, TAS_D) of $AG_{res}=2542, 10242, 40962, 40962$. Prior to

$aMSM_{STR}$ registration midthickness surfaces were first normalised to have constant surface area between timepoints; this encourages even penalisation of shape changes across the surface allowing for regions to deform at different rates. In each case λ was selected so as to optimise feature map correlation relative to areal distortions.

After registrations, true midthickness surfaces were projected through the estimated warp using the procedure described in sec. 4. Methods were then compared in terms of the goodness of fit of the obtained alignments, relative to distortions of the anatomy. In order to compare deformations between data sets relative values for areal ($\log_2 J$) and shape ($\log_2 R$) distortions were estimated at each vertex i as $\hat{Log}_2 J_i = \log_2 \frac{J_i}{\bar{J}}$, $\hat{Log}_2 R_i = \log_2 \frac{R_i}{\bar{R}}$. Alignment quality was assessed through improvements in correlation of curvature feature maps and dice overlap of 16 folding-based cortical regions (Fig. 9) relative to affine registration. Here, cortical labels were obtained through registration of each timepoint's T2 brain volume to 20 manually annotated neonatal atlases (ALBERTS: (Gousias et al., 2012)). The resulting 20 segmentations were then fused in a locally-weighted scheme to form the subject's cortical labels (see LWV-MSD in Artachevarria et al. (2009); Makropoulos et al. (2014)), so that similar patches between each atlas and the image have increased weighting.

Results in Fig. 8 show strong improvements for $aMSM_{STR}$ over the spherical methods. Areal distortions are much reduced (Fig 8 a), and shape distortions are lower than all methods other than SD , for which alignment quality is comparatively reduced (Fig. 9).

To further assess the smoothness and consistency of the distortions be-




Figure 7. Comparison of surface geometry (a), cortical labels (b) and curvature maps (c) for one exemplar data set (top=TP1, bottom=TP2)

tween subjects, the initial time point for all scans was registered to a 34 week surface template (Božek et al., 2016), using $sMSM_{STR}$ alignment of sulcal depth maps using the following parameters: $CP_{res}=162, 642, 2542$; $DP_{res}=2542, 10242, 40962$; $\lambda = 0.5, 0.5, 0.5$; $\sigma_{in} = \sigma_{ref} = 6, 4, 2$. Distributions of $LogJ$ for each subject were then resampled onto the template and compared using FSL’s PALM (Winkler et al., 2014), which performs permutation testing for surface image data. This assesses at each vertex whether distortions are statistically greater than zero.

Results in Fig. 10 show mean \hat{Log}_2J across the surface is much smoother for $aMSM_{STR}$ than for spherical methods ($SD, sMSM_{PAIR}, sMSM_{STR}$). This translates to much broader areas where distortions are significantly above zero. These areas correspond to regions in the frontal and parietal lobe, which are expected to grow faster during this time period as well as after birth (Moeskops et al., 2015; Hill et al., 2010).




Figure 8. Cumulative distribution functions of $\hat{\text{Log}}_2 J$ and $\hat{\text{Log}}_2 R$. for different methods: $sMSP_{PAIR}$ (red); SD (green); $sMSP_{STR}$ (blue); $aMSP_{STR}$ (black). Functions are estimated from the full distribution of strain values estimated by combining per-vertex strain values across all 10 deformations.




Figure 9. Alignment quality of longitudinal warps, assessed through feature map cross correlation and Dice Overlap (averaged across 16 cortical regions). Colour as for Fig. 8

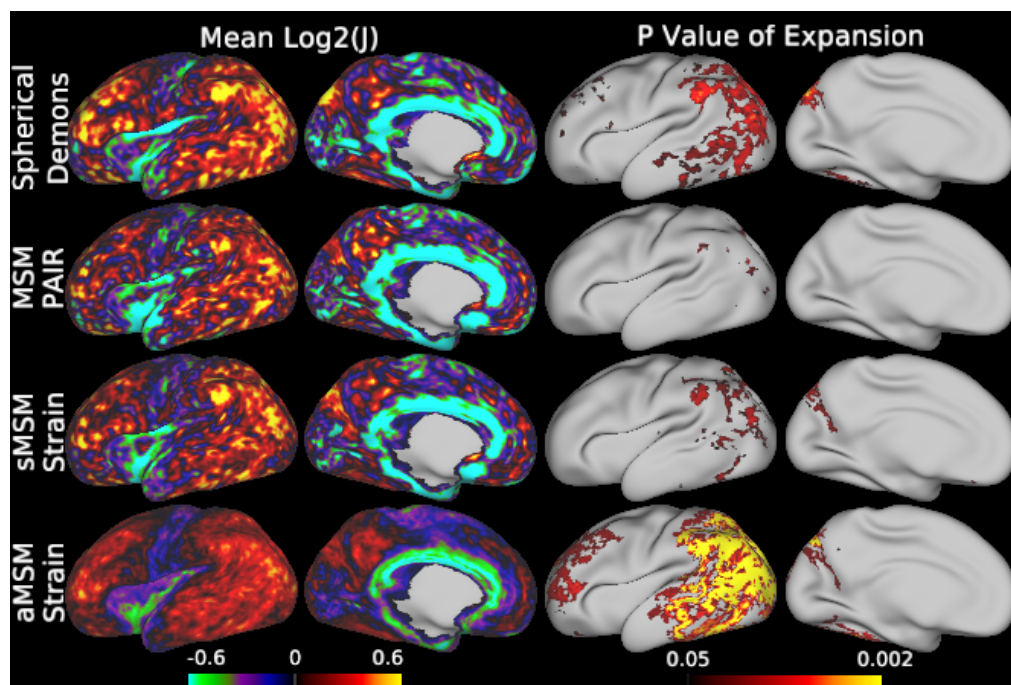


Figure 10. Comparison of distortion fields across subjects. Left: $\hat{L}og_2J$ relative isotropic strain averaged across all 10 neonatal subjects in template space; Right: pvalues for statistical comparison, thresholded at $p < 0.05$

7. Discussion

Human brain imaging studies are extremely diverse; a wide variety of different tissue properties are studied using a range of available imaging modalities. The relationships between these properties are unknown but are likely to be complex since studies have already shown disassociations between patterns of functional and folding organisation (Amunts et al., 2007; Glasser et al., 2016a). Despite this issue the majority of spatial normalisation techniques focus on alignment of specific feature types, including: cortical folding patterns (Fischl et al., 1999b; Yeo et al., 2010; Wright et al.,

2015; Lombaert et al., 2013), cortical geometry (Durrleman et al., 2009), tfMRI time series (Sabuncu et al., 2010), patterns of functional organisation (Conroy et al., 2013; Nenning et al., 2017), and fixed combinations of geometry and fMRI/dMRI (Lombaert et al., 2015; Orasanu et al., 2016b). By contrast in this paper we present MSM, a method which allows flexible and robust alignment of a wide variety of different combinations of features on the cortical surface, as well as biologically-constrained and statistically-plausible deformations of cortical anatomy.

This paper builds on previous work in which we proposed a tool for spherical alignment of brain imaging data, implemented using discrete optimisation (Robinson et al., 2014). This framework offered significant advantages for multimodal registration on account of offering a modular and flexible choice of cost functions, and deformations robust to local minima. One limitation of the original approach was that it was implemented through use of first-order discrete methods and a regularization function with several limitations (see introduction). This resulted in a tension between controlling peak distortions in regions with individual differences in the layout of cortical areas and maximizing alignment across the rest of the surface. This paper therefore presents an improved framework that utilises advances in discrete optimisation that allow for the inclusion of higher-order smoothness penalty terms (Ishikawa, 2009, 2014). Using this framework we improve the robustness of MSM through inclusion of a new hyperelastic strain energy density penalty (Eq. 4) that allows complete control over local changes in shape (R) and area (J). In this way we have been able to significantly improve the alignment of complex multi-modal feature sets. In particular, enhancing the alignment

of data from the adult multimodal parcellation feature set from the HCP Glasser et al. (2016a).

A further advantage of the proposed high-order framework is that it has enabled modification of MSM to allow alignment of cortical anatomies, by imposing regularisation on the anatomical deformations implied by the spherical warp. Experiments performed for longitudinal alignment of 10 neonatal cortices imaged twice between two specific time points: 32.66 ± 1.22 weeks PMA and 41.47 ± 1.61 weeks PMA show that anatomical MSM (*aMSM*) generates smooth and biologically plausible deformations. These reflect patterns of cortical growth similar to those previously reported in region-of-interest studies (Moeskops et al., 2015). But, importantly, *aMSM_{STR}* allowed us to observe these trends using a much smaller sample size, and without the loss of detail associated with large regions of interest.

Throughout this paper we have compared the proposed MSM only against other spherical registration frameworks: Spherical Demons (*SD*) and FreeSurfer (*FS*), and solely optimised for cortical folding alignments. We compare directly to spherical methods as these offer more flexibility in terms of the range of features that can be used to drive the registration. Whilst it is possible that methods, such as Durrleman et al. (2009); Lombaert et al. (2013); Orasanu et al. (2016a), designed purposely for smooth and/or diffeomorphic alignment of cortical surface geometries may outperform MSM for the specific task of longitudinal alignment, these methods are coupled to alignment of cortical shape. This limits their flexibility for between subject alignment on account of known dissociation between the brain’s functional (or areal) organisation and patterns of cortical folding (Amunts et al., 2007; Glasser et al., 2016a;

Nenning et al., 2017). For the same reason, it is not clear how methods that combine spectral alignment of shape with correspondences learnt from functional or diffusion MRI (Lombaert et al., 2015; Orasanu et al., 2016a) should resolve this conflict to obtain a unified mapping across the whole brain.

We compare against SD and FS only for folding alignment on account of the fact that FS registration is fixed and immutable (allowing cortical folding alignment only) and the current implementation of SD allows only for alignment of univariate features. It is true that other groups have applied SD to alignment of brain function by mapping or embedding functional data to a univariate space (Nenning et al., 2017; Tong et al., 2017). However, we have shown through comparisons of SD and $sMSM_{STR}$ for registration of adult folding data that $sMSM_{STR}$ achieves warps with comparable smoothness to SD . By contrast the strength of MSM is that it works flexibly with a range of multivariate and multimodal features as well as overcoming the limitations of traditional spherical methods by generating plausible deformations of cortical surface anatomies. For these reasons we consider an extensive comparison against MSM against these specialised cases outside of the scope of this paper, but would welcome an independent review on these topics.

There remain limitations of the proposed approach inasmuch as the current implementation, using Ishikawa (2009, 2014), reduces the higher-order problem to a series of binary problems. In this paper, a multi-label solution is obtained through use of the fusion moves technique (Lempitsky et al., 2010). However, this circumvents the fast multi-label optimisation offered by the FastPD algorithm (Komodakis and Tziritas, 2007; Komodakis et al., 2008), leading to a slower solution. Specifically, on a 64bit linux system,

the proposed version of MSM runs in approximately 1hr 15 minutes (comparable in run time to FreeSurfer) whereas the original pairwise form runs in 10mins, and Spherical Demons runs in less than 5 mins. Run times can be considerably brought down through appropriate code parallelisation. However, future work should also explore alternative higher-order optimisation strategies such as Fix et al. (2014); Komodakis and Paragios (2009).

Despite the improvement in robustness of the proposed method to the effects of noise and topological variance in the data, this method is still a spatially-smooth, topology-constraining registration approach. An important future avenue will be to address the significant limitation that spatially constrained deformation approaches cannot align brains with variable functional topologies, such as those observed for in $\sim 10\%$ of subjects for area 55b in Glasser et al. (2016a). One avenue may be to explore combining (Langs et al., 2010; Haxby et al., 2011) or graph matching (Ktena et al., 2016) with spatially-constrained registration approaches, such that constraints are placed to ensure that regions cannot be matched if they are very far apart in space (Jordan et al., 2016). Alternatively, in Robinson et al. (2016) we propose a group-wise registration scheme that accounts for topological variation through minimisation of rank of the feature set across the group.

In conclusion, these results underline MSM as a truly flexible tool, which provides a singular resource to study a wide variety of properties of brain organisation across a range of populations. The strain-based version of MSM will be used in HCP studies on development, aging, and disease. Future work will expand studies on longitudinal fetal and neonatal cortical development across a larger cohort to better understand the mechanisms underpin-

ning cortical growth, and will extend studies of neonatal resting-state networks through development of spatio-temporal templates of brain functional and structural organisation Božek et al. (2016). By quantifying patterns of structural and functional development it will be possible to generate vital biomarkers indicating neurodevelopmental outcomes for vulnerable groups such as preterm infants.

8. Acknowledgements

The research leading to these results has received funding from the European Research Council under the European Unions Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 319456. We are grateful to the families who generously supported this trial. The work was supported by the NIHR Biomedical Research Centers at Guys and St Thomas NHS Trust

Data from the Human Connectome Project (HCP - WU-Minn Consortium, Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) was funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. M.G. is supported by an individual fellowship F30-MH097312 (NIH). KG was supported by grant T32 EB018266 (NIH).

Abdollahi, R. O., Kolster, H., Glasser, M. F., Robinson, E. C., Coalson, T. S., Dierker, D., Jenkinson, M., Van Essen, D. C., Orban, G. A., 2014. Correspondences between retinotopic areas and myelin maps in human visual cortex. *Neuroimage* 99, 509–524.

- Amunts, K., Malikovic, A., Mohlberg, H., Schormann, T., Zilles, K., 2000. Brodmann's areas 17 and 18 brought into stereotaxic space where and how variable? *Neuroimage* 11 (1), 66–84.
- Amunts, K., Schleicher, A., Zilles, K., 2007. Cytoarchitecture of the cerebral cortex - more than localization. *NeuroImage* 37 (4), 1061–1065.
- Artaechevarria, X., Munoz-Barrutia, A., Ortiz-de Solórzano, C., 2009. Combination strategies in multi-atlas image segmentation: Application to brain mr data. *IEEE transactions on medical imaging* 28 (8), 1266–1277.
- Barch, D. M., Burgess, G. C., Harms, M. P., Petersen, S. E., Schlaggar, B. L., Corbetta, M., Glasser, M. F., Curtiss, S., Dixit, S., Feldt, C., et al., 2013. Function in the human connectome: Task-fMRI and individual differences in behavior. *NeuroImage* 80, 169–189.
- Boykov, Y., Veksler, O., Zabih, R., 2001. Fast approximate energy minimization via graph cuts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 23 (11), 1222–1239.
- Božek, J., Fitzgibbon, S., Wright, R., Rueckert, D., Jenkinson, M., Robinson, E. C., 2016. Construction of a neonatal cortical surface atlas using multimodal surface matching. In: *2016 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*.
- Conroy, B. R., Singer, B. D., Guntupalli, J. S., Ramadge, P. J., Haxby, J. V., 2013. Inter-subject alignment of human cortical anatomy using functional connectivity. *NeuroImage* 81, 400–411.

- Cordero-Grande, L., Teixeira, R. P. A., Hughes, E. J., Hutter, J., Price, A. N., Hajnal, J., 2016. Sensitivity encoding for aligned multishot magnetic resonance reconstruction. *IEEE Trans. Comput. Imaging* 2, 266–280.
- Durrleman, S., Pennec, X., Trouvé, A., Ayache, N., 2009. Statistical models of sets of curves and surfaces based on currents. *Medical image analysis* 13 (5), 793–808.
- Durrleman, S., Pennec, X., Trouvé, A., Braga, J., Gerig, G., Ayache, N., 2013. Toward a comprehensive framework for the spatiotemporal statistical analysis of longitudinal shape data. *International journal of computer vision* 103 (1), 22–59.
- Fischl, B., Rajendran, N., Busa, E., Augustinack, J., Hinds, O., Yeo, B. T., Mohlberg, H., Amunts, K., Zilles, K., 2008. Cortical folding patterns and predicting cytoarchitecture. *Cerebral Cortex* 18 (8), 1973–1980.
- Fischl, B., Sereno, M. I., Dale, A. M., 1999a. Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage* 9 (2), 195–207.
- Fischl, B., Sereno, M. I., Tootell, R. B., Dale, A. M., et al., 1999b. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human brain mapping* 8 (4), 272–284.
- Fix, A., Wang, C., Zabih, R., 2014. A primal-dual algorithm for higher-order multilabel markov random fields. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1138–1145.

- Frost, M. A., Goebel, R., 2013. Functionally informed cortex based alignment: an integrated approach for whole-cortex macro-anatomical and roi-based functional alignment. *Neuroimage* 83, 1002–1010.
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Jesper, A., Beckmann, C. F., Jenkinson, M., Smith, S. M., Van Essen, D. C., 2016a. A multi-modal parcellation of human cerebral cortex. *Nature* in press.
- Glasser, M. F., Smith, S. M., Marcus, D. S., Andersson, J. L., Auerbach, E. J., Behrens, T. E., Coalson, T. S., Harms, M. P., Jenkinson, M., Moeller, S., et al., 2016b. The human connectome project’s neuroimaging approach. *Nature Neuroscience* 19 (9), 1175–1187.
- Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., et al., 2013. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*.
- Glasser, M. F., Van Essen, D. C., 2011. Mapping human cortical areas in vivo based on myelin content as revealed by T1-and T2-weighted MRI. *The Journal of Neuroscience* 31 (32), 11597–11616.
- Glocker, B., Heibel, T. H., Navab, N., Kohli, P., Rother, C., 2010. Triangleflow: Optical flow with triangulation-based higher-order likelihoods. In: *Computer Vision–ECCV 2010*. Springer, pp. 272–285.
- Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N., 2008.

- Dense image registration through MRFs and efficient linear programming. *Medical image analysis* 12 (6), 731–741.
- Glocker, B., Sotiras, A., Komodakis, N., Paragios, N., 2011. Deformable medical image registration: Setting the state of the art with discrete methods*. *Annual review of biomedical engineering* 13, 219–244.
- Gordon, E. M., Laumann, T. O., Adeyemo, B., Gilmore, A. W., Nelson, S. M., Dosenbach, N. U., Petersen, S. E., 2017. Individual-specific features of brain systems identified with resting state functional correlations. *NeuroImage* 146, 918–939.
- Gousias, I. S., Edwards, A. D., Rutherford, M. A., Counsell, S. J., Hajnal, J. V., Rueckert, D., Hammers, A., 2012. Magnetic resonance imaging of the newborn brain: manual segmentation of labelled atlases in term-born and preterm infants. *Neuroimage* 62 (3), 1499–1509.
- Gu, X., Wang, Y., Chan, T. F., Thompson, P. M., Yau, S.-T., 2004. Genus zero surface conformal mapping and its application to brain surface mapping. *Medical Imaging, IEEE Transactions on* 23 (8), 949–958.
- Harrison, S. J., Woolrich, M. W., Robinson, E. C., Glasser, M. F., Beckmann, C. F., Jenkinson, M., Smith, S. M., 2015. Large-scale probabilistic functional modes from resting state fmri. *Neuroimage* 109, 217–231.
- Haxby, J. V., Guntupalli, J. S., Connolly, A. C., Halchenko, Y. O., Conroy, B. R., Gobbini, M. I., Hanke, M., Ramadge, P. J., 2011. A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron* 72 (2), 404–416.

- Hill, J., Dierker, D., Neil, J., Inder, T., Knutsen, A., Harwell, J., Coalson, T., Van Essen, D., 2010. A surface-based analysis of hemispheric asymmetries and folding of cerebral cortex in term-born human infants. *The Journal of Neuroscience* 30 (6), 2268–2276.
- Hughes, E. J., Winchman, T., Padormo, F., Teixeira, R., Wurie, J., Sharma, M., Fox, M., Hutter, J., Cordero-Grande, L., Price, A. N., et al., 2016. A dedicated neonatal brain imaging system. *Magnetic Resonance in Medicine*.
- Jordan, M. C., Joulin, A., Beck, D. M., Fei-Fei, L., 2016. Locally-optimized inter-subject alignment of functional cortical regions. arXiv preprint arXiv:1606.02349.
- Ishikawa, H., 2009. Higher-order clique reduction in binary graph cut. In: *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, pp. 2993–3000.
- Ishikawa, H., 2014. Higher-order clique reduction without auxiliary variables. In: *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*. IEEE, pp. 1362–1369.
- Knutsen, A. K., Chang, Y. V., Grimm, C. M., Phan, L., Taber, L. A., Bayly, P. V., 2010. A new method to measure cortical growth in the developing brain. *Journal of biomechanical engineering* 132 (10), 101004.
- Komodakis, N., Paragios, N., 2009. Beyond pairwise energies: Efficient optimization for higher-order mrfs. In: *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, pp. 2985–2992.

- Komodakis, N., Tziritas, G., 2007. Approximate labeling via graph cuts based on linear programming. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 29 (8), 1436–1453.
- Komodakis, N., Tziritas, G., Paragios, N., 2008. 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.
- Ktena, S. I., Parisot, S., Passerat-Palmbach, J., Rueckert, D., 2016. Comparison of brain networks with unknown correspondences. In: *22nd Annual Meeting of the Organization for Human Brain Mapping*.
- Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M. A., Hajnal, J. V., Schnabel, J. A., 2012. Reconstruction of fetal brain mri with intensity matching and complete outlier removal. *Medical image analysis* 16 (8), 1550–1564.
- Langs, G., Tie, Y., Rigolo, L., Golby, A., Golland, P., 2010. Functional geometry alignment and localization of brain areas. In: *Advances in neural information processing systems*. pp. 1225–1233.
- Lempitsky, V., Rother, C., Roth, S., Blake, A., 2010. Fusion moves for markov random field optimization. *IEEE transactions on pattern analysis and machine intelligence* 32 (8), 1392–1405.
- Lombaert, H., Arcaro, M., Ayache, N., 2015. Brain transfer: spectral analysis of cortical surfaces and functional maps. In: *International Conference on Information Processing in Medical Imaging*. Springer, pp. 474–487.

- Lombaert, H., Sporring, J., Siddiqi, K., 2013. Diffeomorphic spectral matching of cortical surfaces. In: *Information Processing in Medical Imaging*. Springer, pp. 376–389.
- Lyu, I., Kim, S. H., Seong, J.-K., Yoo, S. W., Evans, A., Shi, Y., Sanchez, M., Niethammer, M., Styner, M. A., 2015. Robust estimation of group-wise cortical correspondence with an application to macaque and human neuroimaging studies. *Frontiers in neuroscience* 9.
- Makropoulos, A., Gousias, I. S., Ledig, C., Aljabar, P., Serag, A., Hajnal, J. V., Counsell, S. J., Rueckert, D., et al., 2014. Automatic whole brain mri segmentation of the developing neonatal brain. *Medical Imaging, IEEE Transactions on* 33 (9), 1818–1831.
- Makropoulos, A., Robinson, E. C., Schuh, A., Wright, R., Fitzgibbon, S., Bozek, J., Counsell, S. J., Steinweg, J., Passerat-Palmbach, J., Lenz, G., et al., 2017. The developing human connectome project: a minimal processing pipeline for neonatal cortical surface reconstruction. *bioRxiv*, 125526.
- Moeller, S., Yacoub, E., Olman, C. A., Auerbach, E., Strupp, J., Harel, N., Uğurbil, K., 2010. Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magnetic Resonance in Medicine* 63 (5), 1144–1153.
- Moeskops, P., Benders, M. J., Kersbergen, K. J., Groenendaal, F., de Vries, L. S., Viergever, M. A., Išgum, I., 2015. Development of cortical morphol-

- ogy evaluated with longitudinal mr brain images of preterm infants. *PloS one* 10 (7), e0131552.
- Neemuchwala, H. F., 2005. Entropic graphs for image registration. Ph.D. thesis, The University of Michigan.
- Nenning, K.-H., Liu, H., Ghosh, S. S., Sabuncu, M. R., Schwartz, E., Langs, G., 2017. Diffeomorphic functional brain surface alignment: Functional demons. *NeuroImage*.
- Orasanu, E., Bazin, P.-L., Melbourne, A., Lorenzi, M., Lombaert, H., Robertson, N. J., Kendall, G., Weiskopf, N., Marlow, N., Ourselin, S., 2016a. Longitudinal analysis of the preterm cortex using multi-modal spectral matching. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 255–263.
- Orasanu, E., Melbourne, A., Cardoso, M. J., Lombaert, H., Kendall, G. S., Robertson, N. J., Marlow, N., Ourselin, S., 2016b. Cortical folding of the preterm brain: a longitudinal analysis of extremely preterm born neonates using spectral matching. *Brain and behavior* 6 (8).
- Robinson, E. C., Glocker, B., Rajchl, M., Rueckert, D., Campus, S. K., 2016. Discrete optimisation for group-wise cortical surface atlasing. In: *7th International Workshop on Biomedical Image Registration* July 1, Las Vegas, Nevada, USA.
- Robinson, E. C., Jbabdi, S., Glasser, M. F., Andersson, J., Burgess, G. C., Harms, M. P., Smith, S. M., Van Essen, D. C., Jenkinson, M., 2014. Msm:

- A new flexible framework for multimodal surface matching. *NeuroImage* 100, 414–426.
- Rother, C., Kolmogorov, V., Lempitsky, V., Szummer, M., 2007. Optimizing binary mrfs via extended roof duality. In: *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on. IEEE*, pp. 1–8.
- Sabuncu, M. R., Singer, B. D., Conroy, B., Bryan, R. E., Ramadge, P. J., Haxby, J. V., 2010. Function-based intersubject alignment of human cortical anatomy. *Cerebral Cortex* 20 (1), 130–140.
- Smith, S. M., Beckmann, C. F., Andersson, J., Auerbach, E. J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D. A., Griffanti, L., Harms, M. P., et al., 2013. Resting-state fMRI in the Human Connectome Project. *NeuroImage* 80, 144–168.
- Tardif, C. L., Schäfer, A., Waehnert, M., Dinse, J., Turner, R., Bazin, P.-L., 2015. Multi-contrast multi-scale surface registration for improved alignment of cortical areas. *Neuroimage* 111, 107–122.
- Tavor, I., Jones, O. P., Mars, R., Smith, S., Behrens, T., Jbabdi, S., 2016. Task-free mri predicts individual differences in brain activity during task performance. *Science* 352 (6282), 216–220.
- Tong, T., Aganj, I., Ge, T., Polimeni, J. R., Fischl, B., 2017. Functional density and edge maps: Characterizing functional architecture in individuals and improving cross-subject registration. *NeuroImage*.
- Tsui, A., Fenton, D., Vuong, P., Hass, J., Koehl, P., Amenta, N., Coeurjolly, D., DeCarli, C., Carmichael, O., 2013. Globally optimal cortical surface

- matching with exact landmark correspondence. In: *Information Processing in Medical Imaging*. Springer, pp. 487–498.
- Van Essen, D. C., 2005. A population-average, landmark-and surface-based (PALS) atlas of human cerebral cortex. *NeuroImage* 28 (3), 635–662.
- Van Essen, D. C., Glasser, M., Dierker, D., Harwell, J., Coalson, T., 2012. Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cerebral Cortex* 22, 2241–2262.
- Wang, D., Buckner, R. L., Fox, M. D., Holt, D. J., Holmes, A. J., Stoecklein, S., Langs, G., Pan, R., Qian, T., Li, K., et al., 2015. Parcellating cortical functional networks in individuals. *Nature neuroscience* 18 (12), 1853.
- Winkler, A. M., Ridgway, G. R., Webster, M. A., Smith, S. M., Nichols, T. E., 2014. Permutation inference for the general linear model. *Neuroimage* 92, 381–397.
- Woolrich, M. W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., Beckmann, C., Jenkinson, M., Smith, S. M., 2009. Bayesian analysis of neuroimaging data in FSL. *NeuroImage* 45 (1), S173–S186.
- Wright, R., Makropoulos, A., Kyriakopoulou, V., Patkee, P., Koch, L., Rutherford, M., Hajnal, J., Rueckert, D., Aljabar, P., 2015. Construction of a fetal spatio-temporal cortical surface atlas from in utero mri: Application of spectral surface matching. *NeuroImage* 120, 467–480.
- Yeo, B. T., Sabuncu, M. R., Vercauteren, T., Ayache, N., Fischl, B., Golland, P., 2010. Spherical demons: fast diffeomorphic landmark-free surface registration. *Medical Imaging, IEEE Transactions on* 29 (3), 650–668.