Construction of a fetal spatio-temporal cortical surface atlas from in utero MRI: Application of spectral surface matching

R. Wright [a,*], A. Makropoulos [a], V. Kyriakopoulou [b], P.A. Patkee [b], L.M. Koch [a], M.A. Rutherford [b], J.V. Hajnal [b], D. Rueckert [a], P. Aljabar [a,c]

[a] Biomedical Image Analysis Group, Department of Computing, Imperial College London, London SW7 2AZ, UK
[b] Centre for the Developing Brain, Division of Imaging Sciences and Biomedical Engineering, King's College London, King's Health Partners, St. Thomas' Hospital, London SE1 7EH, UK
[c] Department of Biomedical Engineering, Division of Imaging Sciences and Biomedical Engineering, King's College London, King's Health Partners, St. Thomas' Hospital, London SE1 7EH, UK

ABSTRACT

In this study, we construct a spatio-temporal surface atlas of the developing cerebral cortex, which is an important tool for analysing and understanding normal and abnormal cortical development. In utero Magnetic Resonance Imaging (MRI) of 80 healthy fetuses was performed, with a gestational age range of 21.7 to 38.9 weeks. Topologically correct cortical surface models were extracted from reconstructed 3D MRI volumes. Accurate correspondences were obtained by applying a joint spectral analysis to cortices for sets of subjects close to a specific age. Sulcal alignment was found to be accurate in comparison to spherical demons, a state of the art registration technique for aligning 2D cortical representations (average Fréchet distance $\approx$ 0.4 mm at 30 weeks). We construct consistent, unbiased average cortical surface templates, for each week of gestation, from age-matched groups of surfaces by applying kernel regression in the spectral domain. These were found to accurately capture the average cortical shape of individuals within the cohort, suggesting a good alignment of cortical geometry. Each spectral embedding and its corresponding cortical surface template provide a dual reference space where cortical geometry is aligned and a vertex-wise morphometric analysis can be undertaken.

© 2015 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Introduction

The cerebral cortex, a thin layer of neural tissue lining the cerebrum, plays a key role in high level brain functions such as memory, sensory perception, language and consciousness (Kandel, 2013). Cortical folding is an important neurodevelopmental process that occurs largely in the third trimester of pregnancy, where the surface of the cerebral cortex transforms rapidly from a smooth sheet into a highly convoluted arrangement of gyri and sulci (Garel et al., 2001). The increase in surface area allows more neurons to be packed into the skull's limited volume and is fundamental to increases in cognitive ability seen in humans and other intelligent species (Lui et al., 2011). It is thought that abnormal cortical folding could lead to neurocognitive deficiencies. For example, Dubois et al. (2008) found that fetuses with intra-uterine growth restriction, which is a risk factor for developing attention deficit hyperactivity disorder and schizophrenia (Geva et al., 2006), had measurable structural abnormalities at birth that were associated with a lower neurobehavioural development score at term equivalent age.

Our previous work focussed on quantifying cortical development, at a global and lobar scale, by computing a single scalar value, summarising the average extent of folding over a given region (Wright et al., 2014).

However, disturbances in development may be more localised and anomalies may remain undetected when averaging over a relatively large region. We therefore seek to characterise folding at any given point on the cortex.

An average model or atlas can facilitate this by providing a densely sampled reference space where the anatomies of a population can be co-aligned and morphometric analysis performed. Magnetic Resonance Imaging (MRI) based atlases have been used previously to align and compare neuroanatomy between cohorts on a voxel-wise basis using voxel-based morphometry (Ashburner & Friston, 2000). This methodology can also be applied to surfaces, undertaking morphometric analysis on a vertex-wise basis using mesh models (Chung et al., 2008). The aim in this paper is to develop an atlas construction framework to generate a set of accurate average cortical surface templates, which will give a longitudinal reference space for surface morphometry. This will enable us to quantify normal and abnormal folding patterns of the cerebral cortex in future work as well as provide a reference space for subsequent work on connectivity analysis in the developing brain.

Constructing a surface atlas requires finding smooth and accurate correspondences between cortical surfaces from different subjects at different ages, which can be a challenging and computationally intensive task. Previous registration based methods for aligning adult cortical surfaces have required some form of preprocessing to produce a 2D surface parameterisation. This is advantageous as a smooth

* Corresponding author.
E-mail address: r.wright11@imperial.ac.uk (R. Wright).

correspondence between cortical surfaces can be estimated more robustly, while reducing the optimisation problem to two dimensions and decreasing the computation required. For example, Van Essen and Drury (1997) flattened the cortical surface by making a series of surface cuts to reduce distortion before warping one flat cortical map to another. Popular methods such as Freesurfer (Fischl et al., 1999), spherical demons (Yeo et al., 2010) and multi-modal surface matching (Robinson et al., 2014), adopt a 2D spherical coordinate system. Through a process of inflation, a surface can be transformed into a sphere while minimising geometric distortion. Surfaces can then be treated as 2D spherical images whose intensity is given by scalar information associated with each vertex (e.g. sulcal depth) to guide the registration.

Spectral graph theory provides a fast alternative to these methods using the eigenvalues and eigenvectors of the graph Laplacian matrix (Chung, 1997). The eigenvalues of the Laplacian may be interpreted as the natural oscillating frequencies, or resonances, of a physical shape. Each of these resonances has a distinct pattern of displacement, a so-called normal mode of the shape, which is given by the eigenvectors. Equivalent locations on similarly shaped surfaces will have a similar relative amplitude and phase, at any given resonant frequency, which allows surface points to be matched in the spectral domain. Accuracy with spectral methods can be limited due to macroscopic variations in brain shape such as expansion or compression that can cause significant changes in a cortical surface's spectral representation. Previous attempts to address this problem include registering the derived spectral representations (Jain & Zhang, 2006; Mateus et al., 2008; Lombaert et al., 2011, 2013a).

More recently, Lombaert et al. (2013b) proposed a novel spectral-based method to acquire a mapping between surfaces, that is comparable in accuracy to Freesurfer (Fischl et al., 1999) and spherical demons (Yeo et al., 2010), while maintaining a considerable speed advantage. Moreover, atlases generated are more consistent when varying the initial reference surface for aligning a group of surfaces. The standard deviation of atlas boundaries was found to be significantly lower for their spectral method (0.0014 mm $\pm$ 0.0009) compared to spherical demons (0.28 mm $\pm$ 0.08). In this method, a preliminary point-wise correspondence between two surface graphs is established from their spectral representations and additional information such as sulcal depth. A novel dual-layered graph is then constructed, by linking the corresponding vertices across the two graphs. The spectral decomposition of this dual-layered graph provides an orthonormal basis where accurate correspondences between the two cortical surfaces can be estimated. Orasanu et al. (2014) used this methodology to analyse cortical development in neonates born very pre-term (around 26 weeks). Cortical surfaces for five subjects were modelled at two time-points, 30 weeks and 40 weeks equivalent gestational ages. For each individual, cortical correspondences were then estimated between the two time-points, enabling longitudinal changes in cortical geometry to be measured at matched locations.

Given the considerable diversity of cortical shape during gestation, a single anatomical atlas may not be appropriate. Indeed, a sulcus on the surface of an older more neurologically developed brain may not yet exist on a younger brain and thus establishing a meaningful correspondence may not be possible. Without any detectable surface sulcation, from some shape descriptor (e.g. mean curvature), correspondences can be determined by the spectral properties of the surfaces alone. Unfortunately, variations in cerebral shape can alter the spectral vibration modes limiting matching accuracy. Furthermore, macroscopic changes in cerebral shape during gestation alter vibration modes over time, therefore accuracy will also fall as the discrepancy in age increases. A spatio-temporal approach (Serag et al., 2012) can be used to avoid these difficulties, constructing an average anatomical template at several time points, within a specific time interval, where anatomy is similar.

To date, there have only been a handful of studies which have constructed atlases of the developing cortex. Hill et al. (2010) constructed a single cortical surface template from neonates born at term to compare hemispheric asymmetries in infants and adults. A landmark based registration algorithm (Van Essen, 2005) was used to align spherical

representations of the cortex. Li et al. (2014) recently constructed an infant spatio-temporal surface atlas, generating atlas templates for 1, 3, 6, 9, 12, 18 and 24 months of age. This work also aligned cortices on the sphere, using spherical demons (Yeo et al., 2010).

A significant portion of the cortical folding process is completed before term, including the formation of primary and secondary sulci (Garel et al., 2001) and disturbances in development during this period may have a profound effect. Preterm neonatal imaging therefore provides an opportunity to visualise an important part of the cortical folding process. However, the complications associated with premature birth mean that this cohort is not ideally suited for characterising normal development (Ajayi-Obe et al., 2000; Kapellou et al., 2006). For this reason, in utero MR imaging may ultimately provide improved visualisation and quantification of the developing cortex over this critical period. Previously, fetal motion has greatly hindered the acquisition of coherent MR images in utero. However, recent advances in image reconstruction (Rousseau et al., 2006; Jiang et al., 2007; Rousseau et al., 2010; Kim et al., 2010; Gholipour et al., 2010; Kuklisova-Murgasova et al., 2012) have allowed the acquisition of high resolution 3D volumes that are free from motion artefacts.

Two studies have constructed average surface templates of the fetal cortex from in utero MRI. Habas et al. (2012) described an image-analysis framework to investigate early folding patterns revealing asymmetric hemispheric development in a cohort of fetuses, from 20 to 28 weeks gestational age. MR image volumes were co-aligned at each week of gestation using image-based registration and an average cortical surface was extracted from the combined subject tissue maps. 3D image-based registration is, however, ill-suited for cortical alignment at older gestational ages, where the geometry of the cortex is significantly more complex.

More recently, Clouchoux et al. (2012) constructed a spatio-temporal surface atlas consisting of four templates, based on twelve in utero MR images of healthy fetuses (25.2 to 35.1 weeks gestational age), enabling a quantitative analysis of cortical development. Surfaces were reconstructed using a deformable model, CLASP (Kim et al., 2005), and correspondences were found on the unit sphere using an iterative group registration scheme (Lyttelton et al., 2007). This methodology was also used to compare development in normal fetuses and fetuses with a severe form of congenital heart disease (Clouchoux et al., 2013), where a delay in cortical development was observed.

Most of the atlas construction techniques discussed have utilised a spherical coordinate system to register cortical surfaces. In this study, we present an alternative spectral method for constructing a spatio-temporal surface atlas of the developing fetal cortex. We generate average surface templates for each week of gestation from 23 to 37 weeks (15 templates in total), from a cohort of 80 healthy fetal subjects. To our knowledge there are no publicly available surface atlases for the developing fetal cortex over this age range built from such a large population, therefore we will make our atlas available to assist future brain development studies.

In the next section, we outline our spatio-temporal framework and template construction methodology. We describe a natural extension to the spectral-based method of Lombaert et al. (2013b) to embed a group of cortical surfaces simultaneously, thus eliminating bias due to the choice of an initial reference surface for alignment. We then demonstrate the use of kernels (Nadaraya, 1964), to compute a spatially and temporally weighted estimate of the average surface boundary, and show how this is readily converted to a mesh representation (Kazhdan & Hoppe, 2013).

We evaluate our methodology by examining sulcal alignment, regional overlaps, and embedding regularity in comparison with spherical demons and conventional 3D surface registration algorithms. We also demonstrate the consistency of the average templates by generating multiple instances from disjoint subsamples of our dataset of varying sizes. The implications of this work for characterising normal and abnormal cortical growth are then discussed alongside the possible directions for future work.
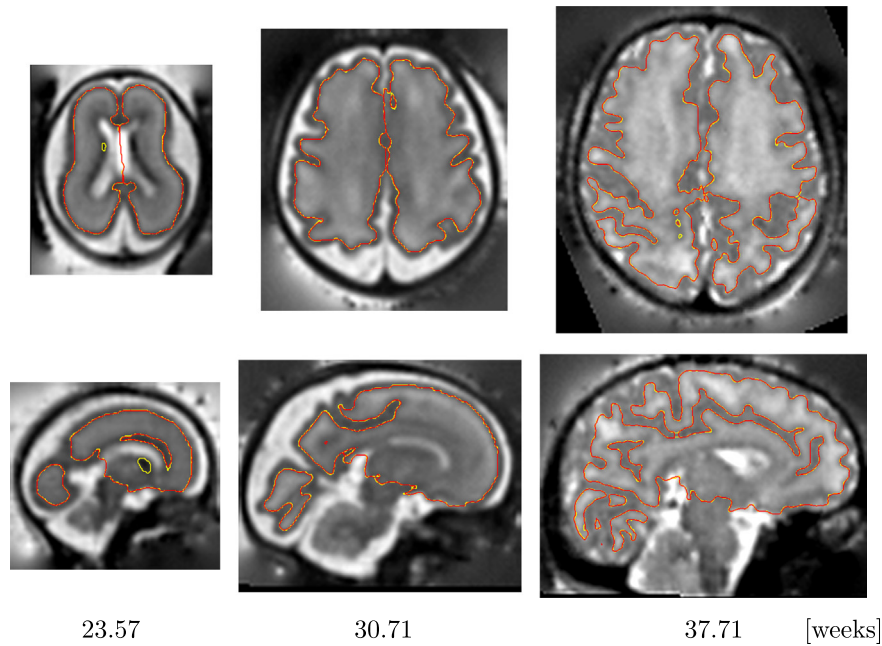
23.57   30.71   37.71   [weeks]

Fig. 1. Cortical surface extraction. Cortical surface models for a selection of gestational ages are overlaid on their corresponding MR image volumes. The isosurface of the sub-cortical tissue probabilities is shown as a yellow contour while the cross-section of the cortical surface model is overlaid as a red contour.

Material and methods

Data and surface modelling

Eighty healthy control subjects with a gestational age (GA) range of 21.7 to 38.9 weeks (mean 29.6 ± 4.6) were used for the study. In utero MR imaging was performed on a 1.5 Philips Intera scanner using a 32 channel coil. Individual slices were acquired using a singe-shot fast spin echo sequence which were then reconstructed into a 3D volume using a robust slice-to-volume registration method (Jiang et al., 2007). Reconstructed volumes had an isotropic voxel resolution of 1 mm. For more details on data preprocessing please see Wright et al. (2014).

As with our previous work (Wright et al., 2014), for the cortical modelling, we focussed on extracting the cortical grey matter (GM)/white matter (WM) surface as this is easier to delineate. MR images were first segmented into basic tissue types using an expectation maximisation (EM) algorithm (Makropoulos et al., 2014), including cortical GM, WM, cerebrospinal fluid (CSF), the deep grey nuclei, lateral ventricles, brain stem and cerebellum. The posteriors obtained from the EM segmentation were then combined to form a probability map for matter lying inside the cortical GM/WM boundary. An explicit surface representation was extracted by deforming a mesh with spherical topology towards the implicit surface, using a similar approach to Kim et al. (2005). The deformation maintains the topology of the mesh and fits the mesh directly to the probability map rather than the image intensities (see Appendix A for more details). All meshes were tested and confirmed to be topologically correct with no self-intersections. To estimate the surface extraction error, a distance map was computed for a segmented boundary (Maurer et al., 2003), and distance values were linearly interpolated at the mesh vertex locations. The average distance error for a vertex was 0.106 mm (±0.006) (see Fig. 1 for surface contour examples).

Framework overview

Given the temporal diversity of anatomy within the cohort, a spatio-temporal framework was adopted, building an average template of cortical shape for each week of gestation across the cohort (Fig. 2). In particular, the approach avoids the difficult task of aligning very disparate cortical anatomies and trying to find accurate correspondences between them. Each template was built from a group of cortical hemisphere surfaces ($N = 19.6 \pm 3.1$) within a narrow time-window of ± 1 week, where cortical features are similar. Both hemisphere surfaces were used to generate the atlas templates (mirroring the right hemisphere), effectively doubling the sample size, allowing a more reliable estimate of the population average surface boundary. By choosing overlapping windows, correspondences between successive embeddings can still be established (Temporal correspondence section). For each template, the contributing surfaces were first affine-aligned by registering all pairs of surfaces, and deforming each surface using the computed average affine transformation. This improved matching accuracy by removing differences in global shape and size. Spectral matching was then performed between all pairs of surfaces, yielding initial vertex-wise correspondences. Surfaces were connected using these preliminary correspondences and a joint spectral analysis gave an embedding in which all the surfaces were co-aligned. The average surface position was then estimated using kernel regression at each point within the embedding, allowing a mean surface to be reconstructed. The following describes these steps in further detail.

Spectral matching

Spectral matching is a basic concept that has been used by several authors (Jain & Zhang, 2006; Mateus et al., 2008; Lombaert et al., 2011, 2013a,b; Orasanu et al., 2014) to find correspondences between two surfaces from their spectra. To recap, we first define a surface model S as a pair $S = \{\mathbf{V}, \mathbf{E}\}$. $\mathbf{V}$ denotes a set of vertices $\mathbf{V} = (\mathbf{x}_1, ..., \mathbf{x}_{|V|})$ with position vectors $\mathbf{x}_i$, while $\mathbf{E}$ denotes an adjacency matrix, where $e_{ij} = 1$, if vertices i and j form an edge, or 0 otherwise. From this we define the $|V| \times |V|$ weighted adjacency matrix $\mathbf{W}$, $w_{ij} = e_{ij} \|\mathbf{x}_i - \mathbf{x}_j\|^{-2}$. The diagonal degree matrix $\mathbf{D} = diag(d_1, d_2 ... d_{|V|})$ is then given by summing all edge weights for each node, $d_i = \sum_j w_{ij}$. The $|V| \times |V|$ general graph Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{W}$ can now be constructed from $\mathbf{W}$ and $\mathbf{D}$.

Solving the eigenvector problem $\mathbf{L} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^{-1}$ gives the diagonal eigenvalue matrix, $\boldsymbol{\Lambda} = diag(\lambda_0, \lambda_1, ..., \lambda_{|V|})$ and the corresponding eigenvector matrix $\mathbf{U} = (\mathbf{u}^0, \mathbf{u}^1, ..., \mathbf{u}^{|V|})$, where each eigenvector $\mathbf{u}^i = (u_{i0}, u_{i2}, ..., u_{i|V|})^T$ is a column of $\mathbf{U}$. Each of the eigenvectors or
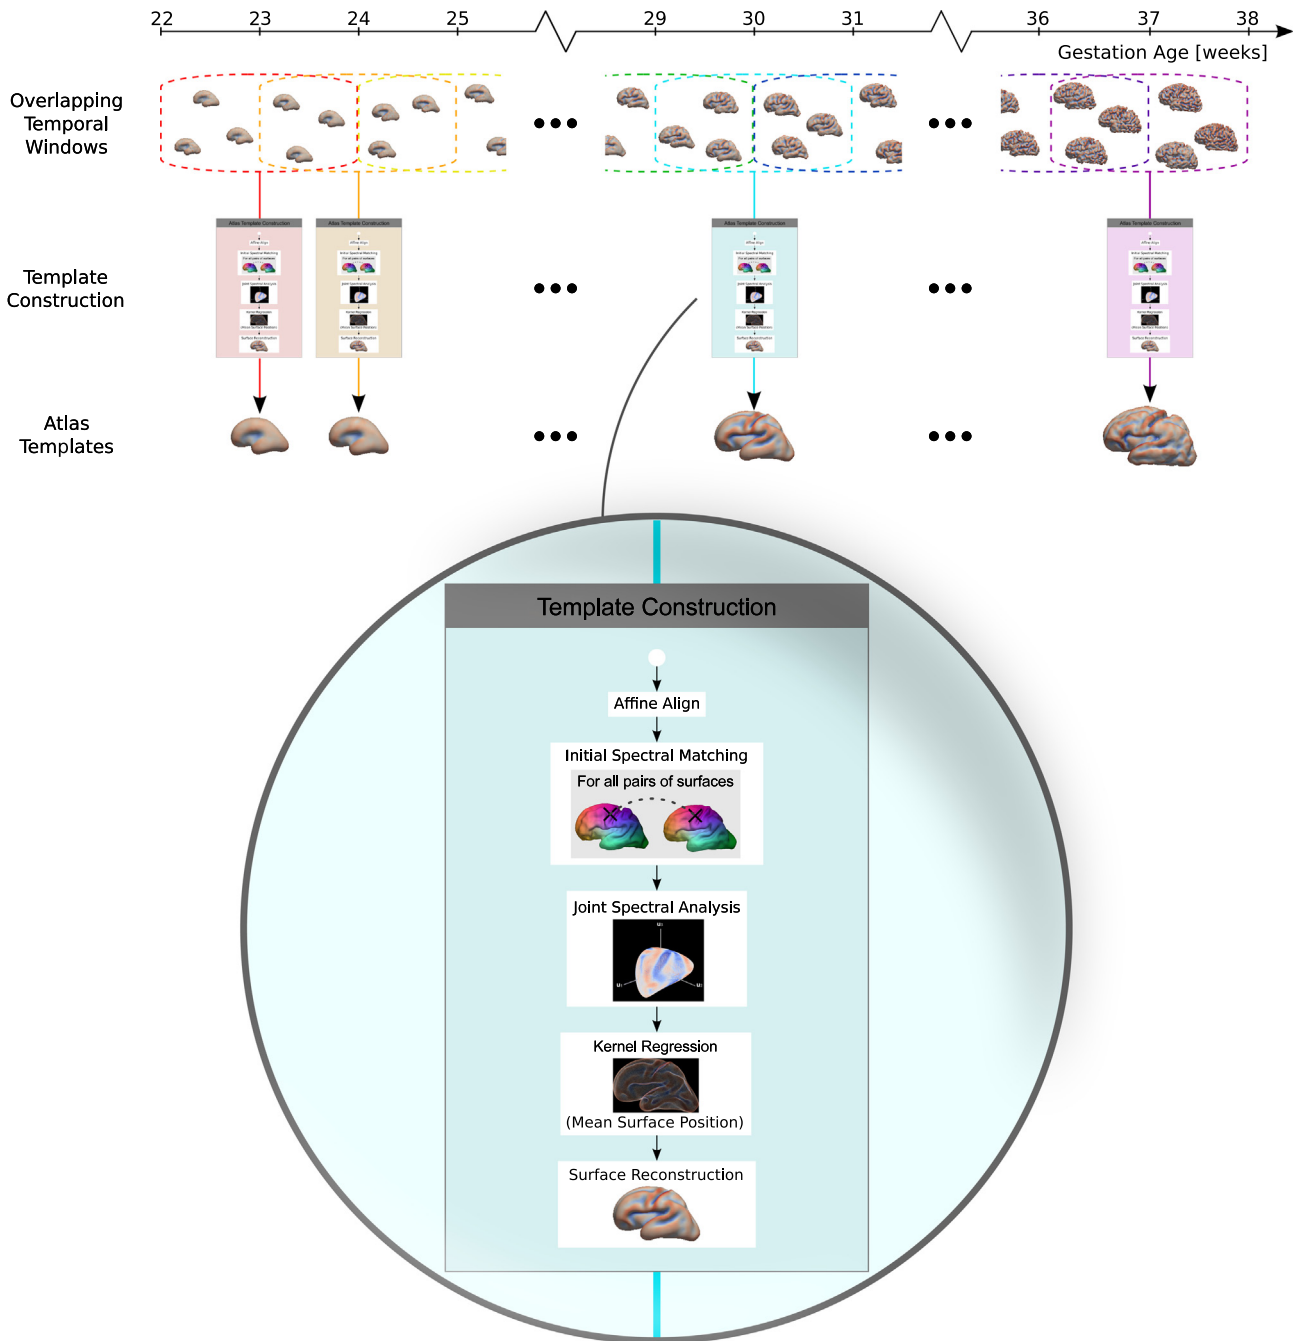
Fig. 2. Framework. An average cortical surface template is constructed for each week of gestation, with all subjects within a week of the target age contributing to the output. A spectral analysis yields spatial correspondences between a group of cortical surfaces, allowing the average surface position to be computed and a template surface constructed. Note the red/blue surface colour mapping depicts the mean curvature of the surfaces.

eigenmodes, $\mathbf{u}^i$, can be viewed as a pattern of vibration for the surface model S (Fig. 3), associated with the resonant frequency λ_i. The first eigenvector is not a vibration mode as $\lambda_0 = 0$ and is thus discarded. These eigenmodes give a useful parameterisation which can be used to match points with another similarly shaped surface. For correct matching, eigenmodes may need to be corrected for their sign ambiguity, multiplicity, and perturbation in isometry (Lombaert et al., 2011, 2013a).

To match points between surfaces the concept of a spectral representation is useful (Fig. 4). The spectral representation of S is denoted as a k-dimensional embedding where the ith vertex has the spectral coordinates $\mathbf{u}'_i = (u_{1i}, u_{2i}, \ldots u_{ki})^T$, which is a row of the truncated matrix $\mathbf{U}_k = (\mathbf{u}^1, \mathbf{u}^2, \ldots, \mathbf{u}^k)$. We define a correspondence map c which gives the

point-wise correspondences between two surfaces S_I and S_J. For each vertex $\mathbf{x}_i \in S_I$ with spectral coordinates $\mathbf{u}'_i$, the corresponding vertex $\mathbf{y}_j \in S_J$ with spectral coordinates $\mathbf{v}'_j$ is determined by the shortest Euclidean distance in the spectral domain, $c(i) = j = \text{argmin}_k \|\mathbf{u}'_i - \mathbf{v}'_k\|$. To achieve a more accurate match, surface descriptors such as mean curvature or the surface normal direction may be added to the embedding as extra dimensions (Lombaert et al., 2011). This ensures that a sulcus (or gyrus) from one surface matches with a nearby sulcus (or gyrus) from another surface. The values for each descriptor may be scaled in order to weigh their influence when finding correspondences.

We have found, however, that this nearest neighbour approach to define an initial correspondence map c between two surfaces
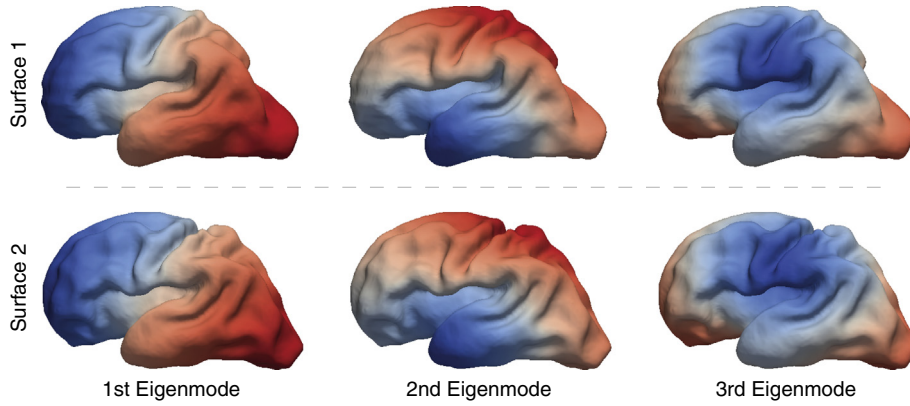
Fig. 3. Surface eigenmodes. The 1st three vibration modes of two surfaces are shown. Each mode depicts a pattern of displacement for a resonant frequency of the surface between poles, colour mapped to blue and red respectively.
Figure adapted from Lombaert et al. (2013b).

(Spectral matching Section) lacks regularisation and the matched locations may be incoherent, i.e. two vertices that are close on one surface may match with vertices that are distant on another surface. This can cause irregularities in the embedded surfaces (Embedding regularity section), particularly when incorporating scalar information into the embedding, which is not determined by surface position (e.g. mean curvature), or where scalar data may be noisy.

To address this, we used edge-based smoothing (Field, 1988) in the spectral domain to regularise the connections between two surfaces S_I and S_J given an initial correspondence map $c = c_{IJ}$ (Algorithm 1). For each vertex $\mathbf{x}_i \in S_I$, with corresponding vertex $\mathbf{y}_{c(i)} \in S_J$ and corresponding spectral coordinate values $\mathbf{v}'_{c(i)}$, we first obtain the smooth spectral coordinate $\bar{\mathbf{v}}'_i$. This is achieved by taking a weighted average of the neighbouring spectral coordinates $\mathbf{v}'_{c(j)}$, based on the edge set $\mathbf{E}_I$, $\bar{\mathbf{v}}'_i = \sum_{j \in N(i)} w_j \mathbf{v}'_{c(j)}$. A more regular correspondence map can then be defined $c(i) = argmin_k \| \bar{\mathbf{v}}'_i - \mathbf{v}'_k \|$.

Algorithm 1. Regularise correspondences

Data: S_I, S_J, correspondence $c = c_{IJ}$
for $i = 1$ **to** $|V_I|$ **do**
$\quad N(i) = \{ j : e_{ij} \in E_I \}$
$\quad \bar{\mathbf{v}}'_i = \sum_{j \in N(i)} w_j \mathbf{v}'_{c(j)}$
$\qquad$ where $\quad w_j \propto 1/\|\mathbf{v}'_{c(i)} - \mathbf{v}'_{c(j)}\|^2 \quad$ and $\quad \sum_{j \in N} w_j = 1$
for $i = 1$ **to** $|V_I|$ **do**
$\quad c(i) = argmin_k \|\bar{\mathbf{v}}'_i - \mathbf{v}'_k\|$

Joint spectral analysis

In Lombaert et al. (2013b), a novel dual-layered surface graph is built from two meshes interconnected by a preliminary correspondence map generated by spectral matching (Spectral matching section). A joint spectral analysis of the connected surfaces produces a set of shared eigenmodes as opposed to two independent sets of eigenmodes. This produces a shared parameterisation of the surfaces.

We extend this approach by building a combined graph for k subjects with surface graphs $S_1, \ldots, S_k$, allowing a joint spectral analysis of all surfaces simultaneously. A multi-layered surface graph $S_c = \{V_c, E_c\}$ is formed where $V_c = (V_1, \ldots, V_k)$ is a concatenation of all mesh vertices and E_c is the combined connectivity matrix. E_c preserves the intra-surface vertex connectivity defined in surface graphs $S_1, \ldots, S_k$ as well as incorporating additional inter-surface connections. For each pair of surfaces S_I and S_J, an edge is formed between each vertex $\mathbf{x}_i \in S^I$ and its corresponding vertex $\mathbf{y}_{c(i)} \in S_J$, where $c = c_{IJ}$ is a preliminary correspondence map between S_I and S_J. As surfaces are linked pairwise, the number of inter-surface edges may be far greater than the number of intra-surface edges. To compensate for this we introduce a parameter ϕ to control the relative weighting of intra- and inter-mesh edges. Intra-mesh edges are weighted as before, by the inverse squared Euclidean distance, $w = \|\mathbf{x}_i - \mathbf{x}_j\|^{-2}$. For inter-surface connections we choose the weighting $w = \frac{3\phi}{k-1}\|\mathbf{x}_i - \mathbf{y}_{c(i)}\|^{-2}$. Note $\phi = 1$ gives an equal weighting between intra-surface and inter-surface edges (see Appendix B for further details).

The spectral decomposition of the graph Laplacian $\mathbf{L}_c = \mathbf{U}_c \mathbf{\Lambda}_c \mathbf{U}_c^{-1}$, built from the multi-layered surface graph S_c, gives a set of shared eigenmodes $\mathbf{U}_c$ and an orthonormal basis where surfaces $S_1, \ldots, S_k$ are co-aligned (Fig. 5).

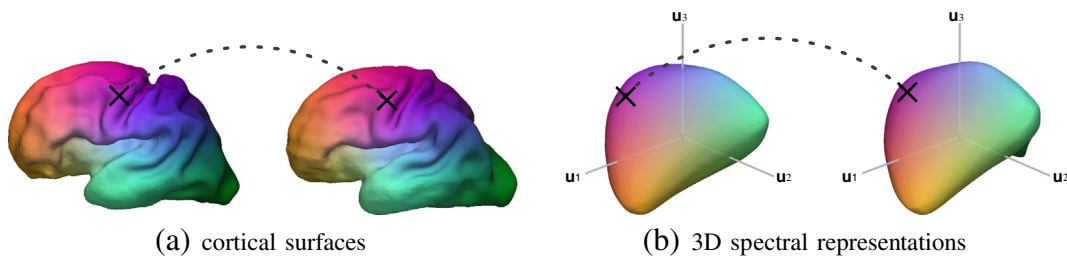(a) cortical surfaces $\qquad\qquad$ (b) 3D spectral representations

Fig. 4. Spectral matching. Vertex-wise correspondences between two surfaces (a) are given by the shortest Euclidean distance in the spectral domain (b). Colour mapping is given by the first three spectral coordinates which are mapped to RGB channels respectively.
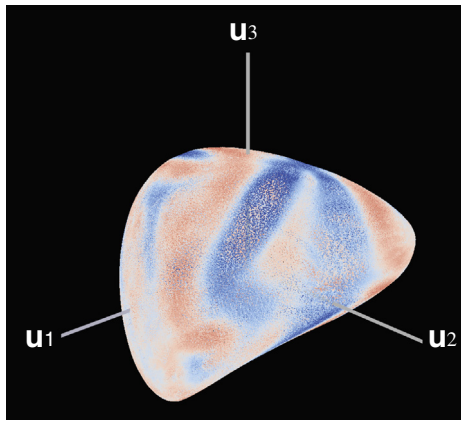Figure adapted from Lombaert et al. (2013b).

Fig. 5. 3D spectral embedding of a group of surfaces. Embedded vertices of 18 cortical surfaces are shown (29-31 weeks GA). Colour mapping depicts the mean curvature at each point on the original surfaces (red, convex; blue, concave).

Kernel regression

The coordinate system provided by the shared eigenmodes (obtained from a joint spectral analysis of a group of surfaces) is referred to as the spectral domain or spectral embedding (Fig. 5). This coordinate system allows us to estimate the values of a scalar variable associated with the surfaces, at matched locations on the cortex. For example, each embedded vertex may have associated shape descriptors (e.g. mean curvature, sulcal depth) or regional anatomical labels (e.g. Brodmann areas) as well as spatial location.

Kernel regression (Nadaraya, 1964) is a non-parametric technique that can be used to estimate the value of one of these scalar variables y, at any given spectral location, by averaging the values of y associated with nearby vertices, using a kernel based weighting function w. Let $\bar{y}_i$ denote the expected value of y at spectral location $\mathbf{u}'_i$ and $N(i)$ a set of vertices close to $\mathbf{u}'_i$. Then $\bar{y}_i$ is given by a weighted average of y for all vertices in $N(i)$, $\bar{y}_i = \sum_{j \in N(i)} y_j w(i,j) / \sum_{j \in N(i)} w(i,j)$. The vertices of each surface are evenly sampled in the spatial domain, however this is not the case in the spectral domain. Therefore, $N(i)$ was defined as the k nearest vertices to $\mathbf{u}'_i$ in order to achieve uniform spatial smoothness.

The weighting function $w(i,j) = w_u\left(\mathbf{u}'_i, \mathbf{u}'_j\right) w_t(t_j) / \mathcal{C}(j)$ was chosen which has three terms: a spatial weighting w_u, a temporal weighting w_t and a count function $\mathcal{C}$. $\mathcal{C}(j)$ simply counts the number of vertices in $N(i)$ that originate from the same surface mesh as j. This removes the influence of different vertex densities for each surface at different locations within the embedding. A Gaussian kernel was used for the spatial weighting $w_u(\mathbf{u}'_i, \mathbf{u}'_j) = \exp\left(-\|\mathbf{u}'_i - \mathbf{u}_j\|^2 / 2\sigma_u^2\right)$ and the temporal weighting $w_t(t_j) = \exp\left(-(t' - t_j)^2 / 2\sigma_t^2\right)$, where t' is a target GA. The temporal smoothness and spatial smoothness of $\bar{y}_i$ are controlled by σ_t and σ_u respectively. σ_u, can be chosen adaptively, so that the average distance of the k vertices is equal to the full width at half maximum of the kernel, $\sigma_u = \frac{1}{2\sqrt{2\ln 2}} \frac{1}{k} \sum_{j \in N(i)} \left\| \mathbf{u}'_i - \mathbf{u}'_j \right\|$. Note, the parameter k now effectively controls the spatial smoothness. Temporal smoothness, σ_t, can be fixed for all time-points or adjusted within regions where the temporal sampling is lower (Serag et al., 2012). This will sacrifice temporal sharpness, when estimating $\bar{y}_i$, in order to smooth out variations in cortical anatomy by incorporating more subjects.

Surface reconstruction

For each embedded vertex i (Fig. 5), the average cortical surface position, $\bar{\mathbf{x}}_i$, can be estimated using kernel regression, given its spectral location $\mathbf{u}'_i$ and a target gestational age t'. This is achieved by treating each spatial dimension as a separate scalar field. The estimated location $\bar{\mathbf{x}}_i$, for each vertex, gives a smooth point sampling of the average surface (Fig. 6). Reconstructing surfaces from point data is a well-studied problem driven by the need for measurement and visualisation of real world objects that can be geometrically sampled using, for example, 3D range scanning technology. We used the popular screened Poisson surface reconstruction algorithm (Kazhdan & Hoppe, 2013), to reconstruct an average surface template from $\bar{\mathbf{x}}$ (Fig. 6). This technique also requires an estimate of the average surface normal direction, $\bar{\mathbf{N}}_i$, for each point sample of the average surface boundary, $\bar{\mathbf{x}}_i$. For further details on how this may be obtained using kernel regression see Appendix C.

Temporal correspondence

Establishing temporal correspondences between embeddings is a core aim of our work as it permits a longitudinal analysis of cortical growth. As a spectral analysis was performed across several overlapping time windows (Fig. 2), for each resulting embedding, a number of the contributing surfaces will also contribute to a neighbouring embedding. For the surfaces that are not shared, we can estimate the position of each of their vertices in the neighbouring embedding. This is achieved by treating the neighbouring coordinates as a labelling associated with each surface vertex within the current embedding. The neighbouring coordinates can then be estimated at unlabelled vertex locations using kernel regression (Kernel regression section), by restricting the k nearest vertices to the set of vertices already labelled.

This procedure is the equivalent of multi-atlas label propagation (Rohlfing et al., 2003; Heckemann et al., 2006), adapted for a set of embedding coordinates. This can also be iterated to propagate a set of coordinates to every surface within the cohort (Fig. 7). This ultimately
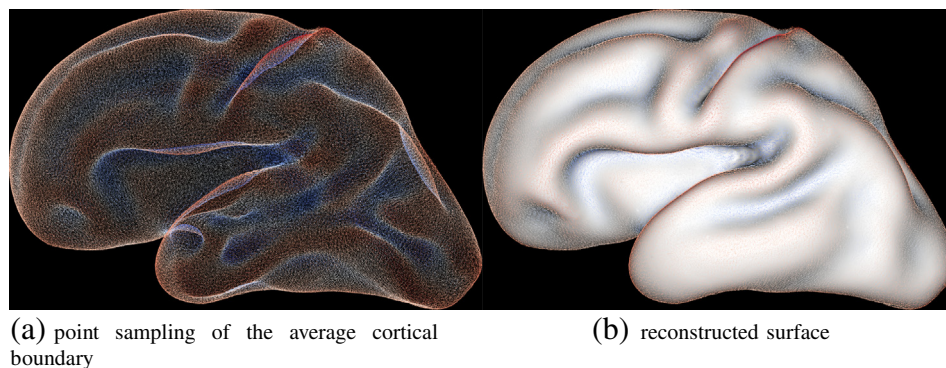

(a) point sampling of the average cortical boundary

(b) reconstructed surface

Fig. 6. Average cortical surface template for a 30 week fetus. Kernel regression in the spectral domain is used to find the average surface position. This gives a point sampling of the average cortical surface (a). Poisson surface reconstruction is then used to extract a closed surface (b), from (a).
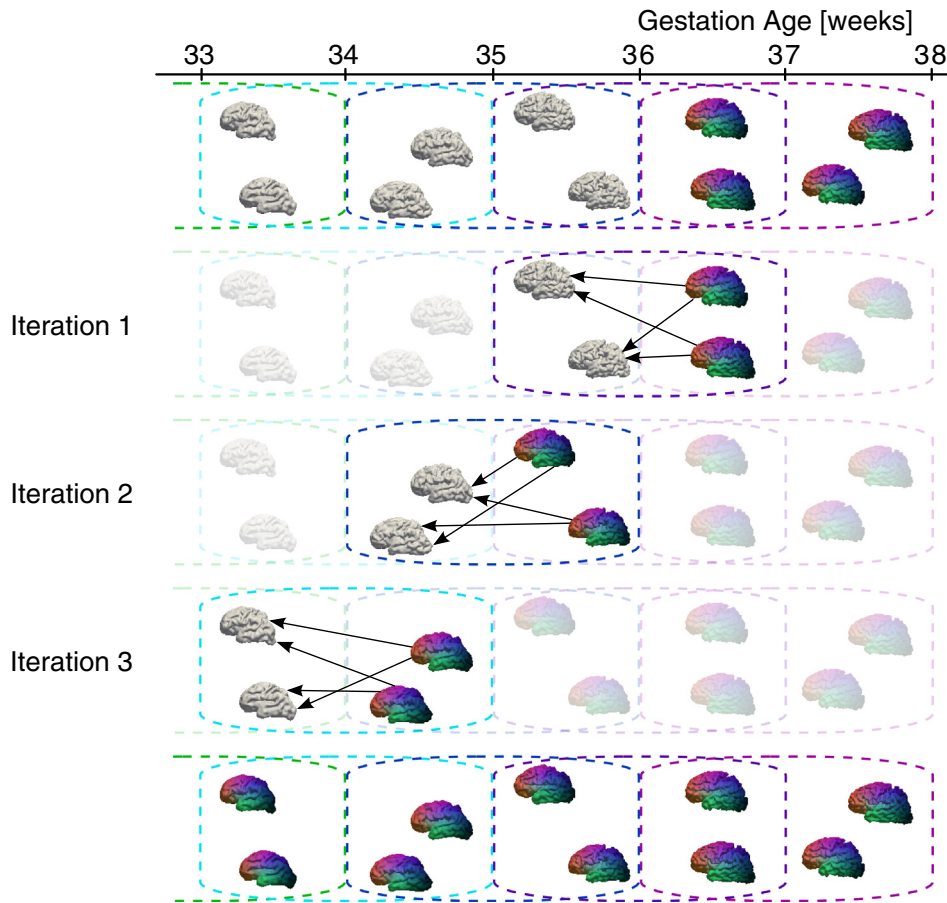
Fig. 7. Stepwise embedding coordinate propagation. This diagram illustrates the propagation of a set of embedding coordinates ($t' = 37$) to the cortical surfaces of younger subjects. Surface colour mapping shows the first 3 embedding coordinates encoded as RGB values. These coordinates can be viewed as a surface labelling, which may be propagated to neighbouring surfaces via surfaces that are shared across neighbouring embeddings. At each iteration, labelled subjects are treated as atlases, whose labels are propagated to any unlabelled surfaces within the same embedding using kernel regression. This process can be repeated iteratively until all surfaces are labelled, yielding a shared parameterisation for all surfaces. The dashed lines enclose the surfaces used to create each of the embeddings.

enables all surfaces to be embedded in a single coordinate system and for correspondences to be found between subjects of different ages even if these differ significantly in their cortical surface geometry (Fig. 8b). We further demonstrate the propagation of nineteen anatomical labels derived from a neonatal atlas set (Gousias et al., 2012). The MRI volumes for fetuses whose age was greater than 36 weeks GA were first parcellated by registering the neonatal atlas set to each volume and fusing the atlas labels. Each surface voxel was
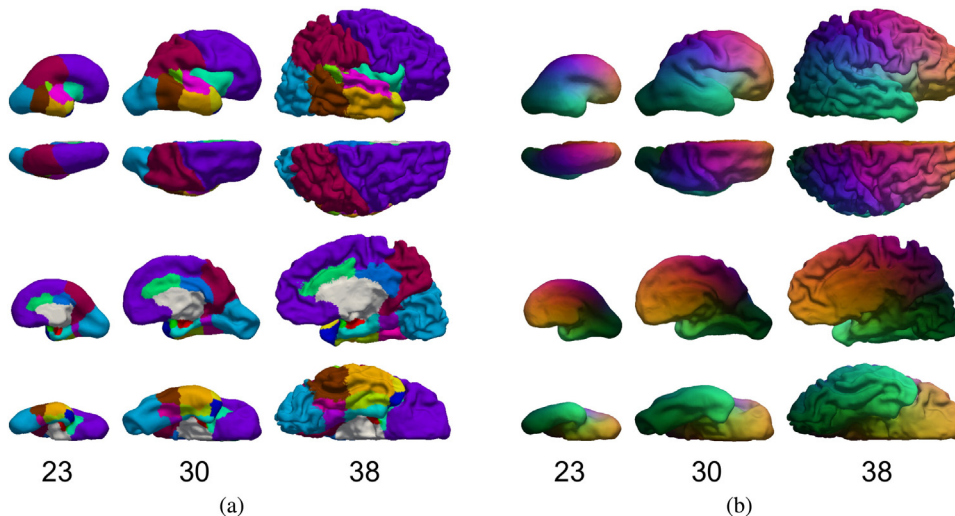

Fig. 8. Multi-atlas label propagation. Three examples of surfaces automatically labelled using kernel regression. (a) A 19 region parcellation. (b) The first embedding coordinates (for $t' = 37$) mapped to RGB channels respectively (see Fig. 7).

then assigned the anatomical label of its nearest voxel from the corresponding MR volume. This formed a surface atlas set for stepwise propagation to all other cortical surfaces (Fig. 8a).

Results and discussion

Model parameters

Graph Laplacian edge weighting: $\phi = 0.1$

The parameter ϕ was introduced to control the weighting of inter-mesh edges relative to intra-mesh edges, when constructing the graph
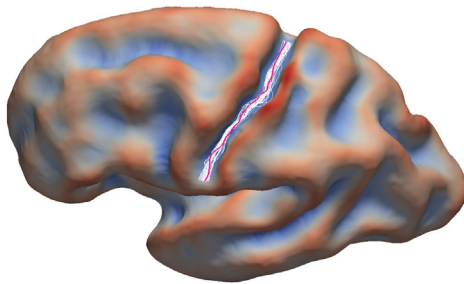

Fig. 9. Aligned central sulci. An example of a manual delineation of the central sulcus (depicted in red), along with 38 central sulcus delineations mapped from other subjects, using the established spectral correspondence (shown in white).

Laplacian for a multi-layered surface graph ($\phi = 1$, gives an equal total weighting to both types of edges). The value of phi was chosen empirically to achieve a compromise between embedding regularity and sulcal alignment accuracy, determined using the validation methods discussed later in this section, and also by visual inspection of the 3D representations of the final embeddings. This value could be tuned further to increase sulcal alignment however, increasing the value of ϕ substantially can have a detrimental effect on embedding regularity.

Kernel parameters: $\sigma_t = 1, \quad k = 100$

The kernel parameters were chosen by observing the smoothness of the final atlas templates. As the subjects in our particular dataset were fairly evenly distributed with age (Wright et al., 2014) a constant temporal smoothness value ($\sigma_t = 1$) was chosen. Both parameters are dependant on the size of the dataset. Additionally the value chosen for k will depend on the required resolution of the surface meshes.

Sulcal alignment accuracy

To determine the precision of the sulcal alignment we manually delineated a common anatomical feature for each subject; the central sulcus. A "polyline" (a line defined by a successive sequence of vertices) was extracted by traversing each surface along the centre of the sulcus. These lines were then mapped pairwise between subjects within each embedding, using the established spectral correspondence (Fig. 9). The similarity between the central sulci after alignment was then assessed using two established measures, the discrete Fréchet distance (F_d) (Alt & Godau,
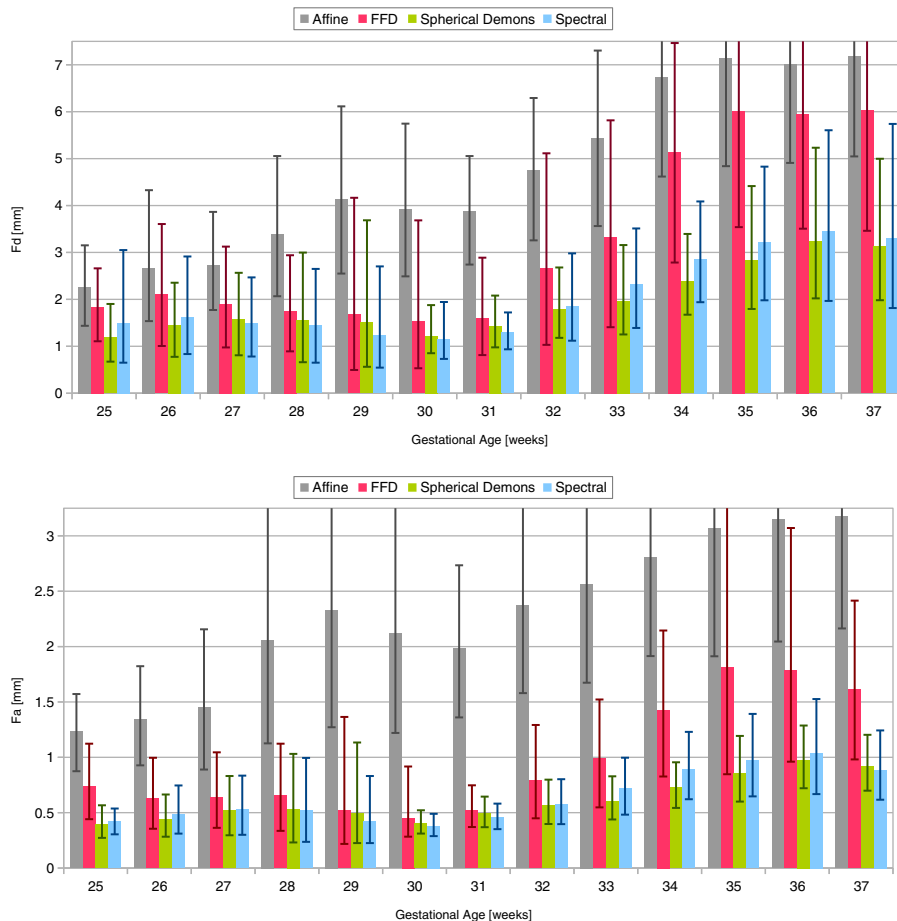

Fig. 10. Intra-embedding sulcal mapping error. For each embedding, the sulcal delineations were mapped pairwise for all surfaces. The average alignment error is shown for four methods, quantified by the Fréchet distance, F_d (top), and the average Fréchet distance, F_a (bottom). Error bars show the standard deviation of the alignment error.

(a) Source gestational age : 27 weeks   (b) Source gestational age : 37 weeks
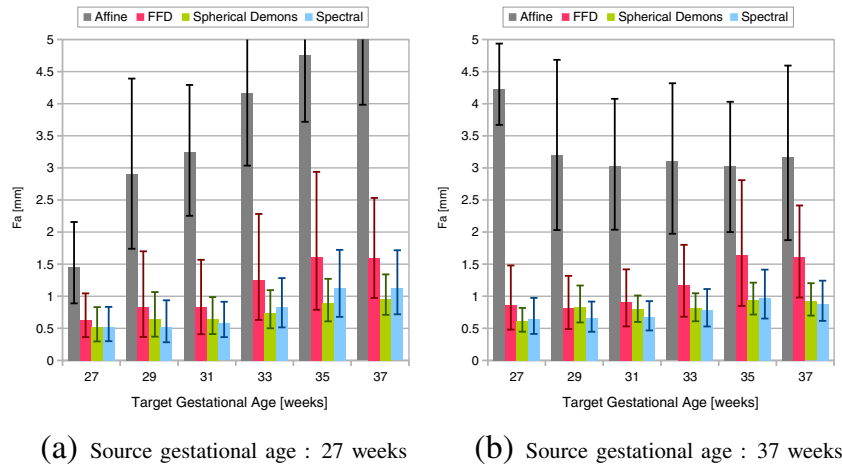
Fig. 11. Inter-embedding sulcal mapping error. Central sulcus delineations were mapped between surfaces that contributed to different embeddings by first establishing temporal correspondences (Temporal correspondence section). The average alignment error, quantified by the average Fréchet distance, F_a, is shown for four methods (a) when mapping delineations for younger source subjects (≈ 27 weeks) to progressively older target subjects and (b) when mapping delineations for older source subjects (≈ 37 weeks) to progressively younger target subjects. Error bars show the standard deviation of the alignment error.

1995) and the average Fréchet distance (F_a) (Brakatsoulas et al., 2005). The central sulcus was selected as it is a well-defined feature that can be easily delineated. It is one of the first sulci to form, and is therefore detectable at an early GA (25 weeks) (Garel et al., 2001) and was observed in most of the subjects within the cohort (>75%).

The Fréchet distance was chosen as it takes into account the location and ordering of the points along a curve, unlike the popular Hausdorff distance, and is therefore more sensitive to discontinuities in mapped polylines. Furthermore, the Fréchet distance has a well defined average (Brakatsoulas et al., 2005) while the Hausdorff distance is a maximum distance measure, increasing its sensitivity to outliers. The Fréchet distance is often referred to as the "dog walker's distance" as it can be described intuitively using the scenario of an owner walking its dog. It is given by the minimum leash length needed to connect a dog and its owner as they traverse two separate paths. The owner and dog may proceed at any speed and even stop, but may not backtrack. Formally, let S be a metric space, then a curve A in S is a continuous mapping from the unit interval $A: [0, 1] \rightarrow S$. A reparameterisation α is a continuous monotonic mapping $\alpha: [0, 1] \rightarrow [0, 1]$. If A and B are two curves in S, then their Fréchet distance $F(A, B)$ is defined as the infimum, over all reparameterisations α and β, of the maximum distance in S between $A(\alpha(t))$ and $B(\beta(t))$, for all $t \in [0, 1]$.

$$F(A, B) = \inf_{\alpha, \beta} \; \max_{t \in [0,1]} \{ d(A(\alpha(t)), B(\beta(t))) \}$$

Here d is a distance function of S.

The sulcal alignment error was compared with that given by spherical demons (Yeo et al., 2010) and two further baseline methods which register surfaces using transformation in 3D: iterative closest point (ICP) registration using affine transformations and ICP using free-form deformations (with a control point spacing of 2.5 mm) (Rueckert et al., 1999).

Alignment error was first compared within each embedding, i.e. between subjects of a similar age. The error between spectral matching and spherical demons was similar, with both outperforming the baseline methods, especially at older ages. The spectral alignment accuracy improved from 25 weeks onwards with the lowest mean error, $F_a \approx 0.4$ mm, at 30 weeks (Fig. 10). This trend was replicated with spherical demons and may be explained by the central sulcus becoming more distinct as it transforms from a shallow depression into a deep sulcus, which is more easily matched. After this the mean matching accuracy dropped for both methods as the cortical complexity increased, with the largest average error seen around 36 weeks ($F_a \approx 1$ mm).

The sulcal mapping accuracy was also compared between subjects of different ages, by first embedding both of their surfaces in the same coordinate system (Temporal correspondence section). Accuracy was then compared when mapping sulcal delineations for younger subjects (≈ 27 weeks) to progressively older subjects (Fig. 11a) and also when mapping sulci for older subjects (≈ 37 weeks) to progressively younger ones (Fig. 11b). Again performance of the two methods was similar and superior to the baseline methods.

Regional overlaps

To assess the alignment accuracy over the whole brain surface, we also evaluated the overlap between different regions of the brain. Each surface was parcellated using a publicly available neonatal MR atlas database (Gousias et al., 2012) into 6 regions (Fig. 12). MR atlases were first registered to each subject's MRI volume using a free-form deformation model (Rueckert et al., 1999), and the atlas labels were fused to create a volumetric parcellation for the subject. This parcellation was then refined using the method of Makropoulos et al. (2014). Each surface vertex was then assigned the anatomical label of its nearest voxel from the corresponding MR volume. Labels were propagated between each pair of surfaces within an embedding using
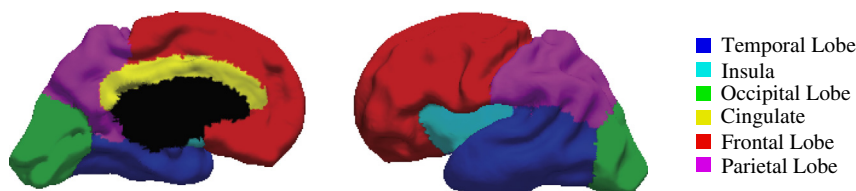

Fig. 12. Surface parcellation. An example of a cortical surface model parcellated into 6 regions.

the spectral (or spherical) correspondences, and their overlap measured using the Dice coefficient (Sørensen, 1948).

The results were again similar between the spectral method and spherical demons, over all six regions (Table 1). As expected, higher Dice scores were observed in the lobar regions that have a larger surface area. The lowest Dice scores were observed for the insula, which is likely due to the absence of a clearly defined inferior anatomical boundary.

Table 1
Cortical region overlap. Region labels were propagated between pairs of cortical surfaces using the established spectral or spherical correspondences, before measuring their overlap using the Dice coefficient. Average Dice scores are summarised.

Region	Spectral	Spherical demons
Temporal lobe	0.937 ($\pm$ 0.015)	0.931 ($\pm$ 0.015)
Insula	0.747 ($\pm$ 0.088)	0.746 ($\pm$ 0.084)
Occipital lobe	0.929 ($\pm$ 0.024)	0.914 ($\pm$ 0.025)
Cingulate	0.829 ($\pm$ 0.045)	0.823 ($\pm$ 0.050)
Frontal lobe	0.964 ($\pm$ 0.009)	0.962 ($\pm$ 0.011)
Parietal lobe	0.922 ($\pm$ 0.020)	0.913 ($\pm$ 0.025)

Table 2
Embedding regularity. Percentage of mesh vertices where structure is preserved (see text).

Measure	Spectral matching		Spherical demons
	Irregular	Regularised	
Percentage	98.35($\pm$1.12)	99.91($\pm$0.09)	99.57($\pm$0.34)
Mean [mm]	0.65($\pm$0.14)	0.40($\pm$0.03)	0.41($\pm$0.032)
Max [mm]	4.15($\pm$2.55)	0.72($\pm$0.23)	0.95($\pm$0.38)

Embedding regularity

The nth eigenvector of the Laplacian has the property of being smooth and monotonous over a surface between at most n poles (Lombaert et al., 2013b). This gives a diffeomorphic mapping between a surface and its spectral representation. The local structure of a surface graph is preserved in the embedding, i.e. an embedded vertex is central to its connected vertices. We conducted a simple experiment to ascertain whether the local structure of each surface graph is preserved when a group of surfaces is embedded using our methodology.

For each vertex i from a surface graph with spatial location $\mathbf{x}_i$ and embedded position $\mathbf{u}'_i$, we use edge-based smoothing (Field, 1988) to find the spectral centroid, $\bar{\mathbf{u}}'_i$, of its neighbouring vertices j, i.e. $\bar{\mathbf{u}}'_i = \sum_{j \in N(i)} w_j \mathbf{u}'_j$, where $w_j \propto 1/\|\mathbf{x}_i - \mathbf{x}_j\|$ and $\sum_{j \in N(i)} w_j = 1$. If the local structure is preserved $\mathbf{u}'_i$ and $\bar{\mathbf{u}}'_i$ should be close. We find the closest vertex to $\bar{\mathbf{u}}'_i$ in the spectral domain, with spatial position $\bar{\mathbf{x}}_i$, and measure the distance $\|\bar{\mathbf{x}}_i - \mathbf{x}_i\|$. This allows us to quantify, for each vertex, how distant its embedded position is from those of its mesh neighbours.

This experiment was repeated for spectral matching with and without regularisation of the surface links (Spectral matching section, final paragraph) and also for spherical demons replacing the spectral embedding coordinates with the method's associated spherical coordinates. For each graph the percentage of voxels where $\|\bar{\mathbf{x}}_i - \mathbf{x}_i\| = 0$ was computed, i.e. the percentage of voxels where the local structure of the graph was preserved. The average percentage was high for spherical demons (99.57%) and both variations of spectral matching (with regularisation, 99.91%, and without, 98.83%) (Table 2). For points where $\|\bar{\mathbf{x}}_i - \mathbf{x}_i\| > 0$, the mean distance was also computed and this was found to be, on average, around the mesh resolution for all methods (mesh edge length ≈ 0.5 mm). The average maximum distance was
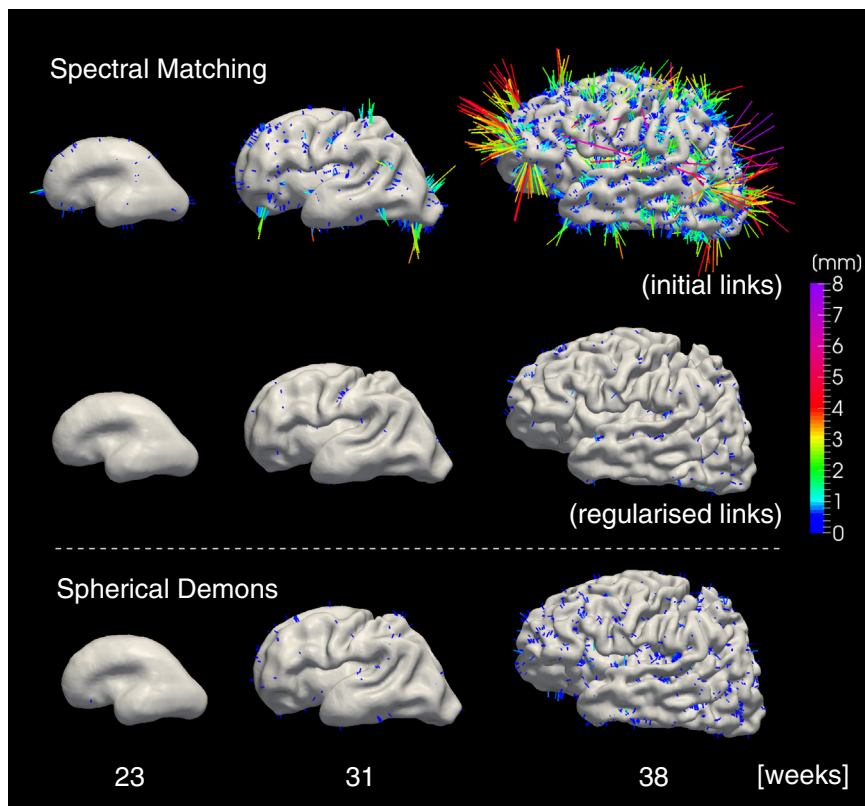

Fig. 13. Embedding regularity. Each line signifies a vertex where the local structure of the surface was not preserved in the embedding. The length of the line is proportional to the distance that the embedded vertex is from the centroid of its neighbouring vertices, when both are mapped back to the spatial domain (the exact distance in mm is given by the colour mapping). Note the irregularities seen when the initial surface links are not regularised.

relatively high for spectral matching with irregular surface links (4.15 mm), however this decreased substantially with regularisation (0.724 mm) in line with spherical demons (0.949 mm).

These results show that the local structure is largely preserved for all methods and that most irregularities are small. However, there were a small number of large irregularities with spectral matching without regularisation, particularly at older gestational ages as the cortical surfaces become more complex (Fig. 13).

Average cortical surfaces

Average cortical surfaces were generated for each week of gestation from the spectral embeddings as well as via spherical demons (Fig. 14). Generating an average surface from a spectral embedding was discussed in Surface reconstruction section, and for spherical demons we follow a similar approach. A regularly sampled spherical mesh gives the correct topology and provides the vertex connectivity for an average surface. This was readily generated by subdividing a platonic solid mesh and projecting the vertices back onto the sphere. Kernel regression (Kernel regression section) can then be used to find an average surface position at each vertex. Transforming the spherical mesh back to the spatial domain, using the average surface position at each vertex, gives an average cortical surface.

The average surfaces for both methods accurately captured the sulcal development of the individuals within the cohort, signifying a good alignment between surfaces. Both techniques produced similar average cortical surfaces, with spherical demons producing slightly sharper surfaces, suggesting a finer sulcal alignment. However, this may also be explained by the slight difference in the way surfaces are reconstructed across the methods. The timing of the sulcation observed in the average surfaces matched previous observations from MR images (Garel et al., 2001). At 23 weeks, the global shape of the brain has already
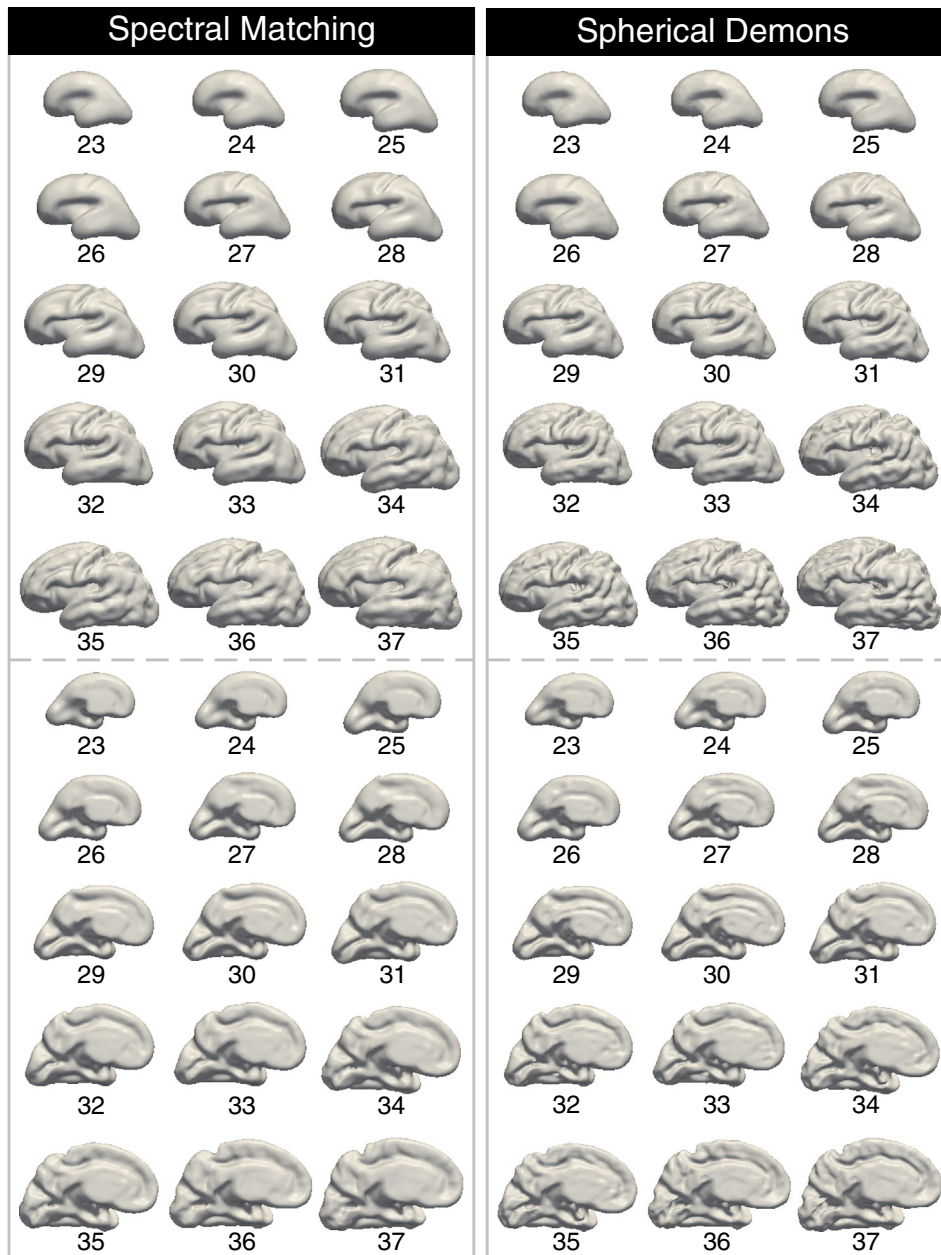

Fig. 14. Average cortical surface templates. Cortical surface templates were constructed for each week of gestation, for both spectral embeddings and spherical demons. Both methods produced visually similar templates, capturing an estimate of the average growth for the cohort.
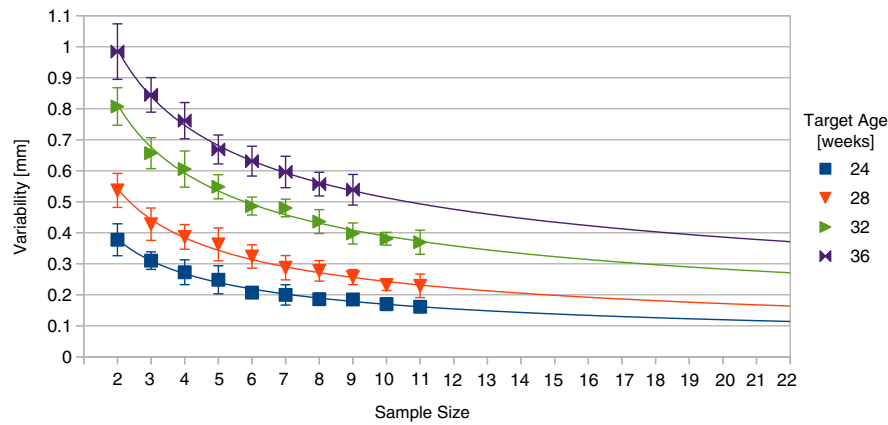
Fig. 15. Effect of sample size on the variability of generated atlas templates. The average distance between distinct templates generated from disjoint subsets of the cortices was computed to give an estimate of the variability of generated surface templates. The points plotted show an average of 10 iterations for a particular sample size and age, with the error bars showing the standard deviation of the 10 iterations. The lines show a power law fit of the data for each target age.

formed and the sylvian, calcarine, parieto-occipital and hipocampic fissures as well as the cingulate are all well defined on the average surfaces. The first sulcus to form is the central sulcus, which is detectable from 24 weeks and very prominent after 27 weeks. By 29 weeks the precentral and postcentral sulci, frontal sulci, superior temporal sulcus and intraparietal sulcus are all clearly visible on the cortical surfaces. Visually the atlas templates are similar to those previously constructed by Habas et al. (2012) and Clouchoux et al. (2013).

Variability of average cortical surfaces

To estimate the variability of generated average surface templates, given a sample size N, we compute two distinct templates from two disjoint subsets, both of size N, of the cortical surface models within a temporal window, and then measure the average distance between them. This was performed for a range of sample sizes (2–11) and 4 target ages (24, 28, 32 and 36 weeks). For each sample size and target age, 10 iterations were performed using randomly generated disjoint pairs of subsets from the data, the average of these 10 iterations is shown as a point in Fig. 15.

At 36 weeks, where anatomical variation is the largest (for the target ages examined), the average surface generated was consistent to around 0.54 mm ± 0.04, for $N = 9$. The relationship between sample size and the average distance between generated average surfaces is well-modelled by a power law curve ($R^2 = 0.992 \pm 0.003$), for sample sizes up to 11. Extrapolating this relationship to sample sizes above 11, we predict the consistency of the average surfaces generated from our complete dataset (Table 3). At 36 weeks, for our full sample size ($N = 18$) we therefore estimate the atlas templates generated to be consistent to around 0.40 mm, with smaller values for ages less than 36 weeks.

Conclusions

In this work, we have constructed a spatio-temporal surface atlas of the developing fetal brain. Establishing a set of preliminary

correspondences between cortical surfaces allowed a group of surface graphs to be linked. This permitted an analysis of their graph Laplacian spectrum as a single entity, yielding an orthonormal basis for modes of an entire group of surfaces. This gave accurate and unbiased correspondences between surfaces, in a fraction of the time required by spherical registration methods.

The sulcal alignment was found to be significantly better than 3D deformation methods and comparable to spherical demons, a state of the art 2D surface registration technique. The regularity of the embedded surfaces was also examined, which is necessary for a diffeomorphic mapping across surfaces, and this was also found to be comparable with that provided by spherical demons.

We have constructed average cortical surface templates for each week of gestation by averaging the surface position over the spectral embeddings. These accurately capture the global shape and the timing of sulcal development over the cortex. Furthermore, the prominence of visible sulcation suggested a good alignment between embedded surfaces. We have also shown the average surfaces generated by our method are reproducible, with an average boundary variability of around 0.54 mm, for disjoint subsamples of our dataset ($N = 9$) for a target age of 36 weeks, where sulcal variation is high.

A spatio-temporal atlas is an important tool for analysing normal and abnormal folding patterns. This work enables a vertex-wise statistical analysis, of morphology across an embedding, which will in turn allow us to identify abnormalities associated with poor neurodevelopmental outcome in the future.

Groups of particular interest include fetal subjects with mild ventricular enlargement, one of the most common congenital anomalies, with a prevalence of around 1%. In most cases, prognosis is good (Gaglioti et al., 2005), however it is associated with a greater risk of cognitive, language, and behavioural impairments in infancy. It is thought that disturbances in cortical development may be indicative of poor outcome in these babies (Kyriakopoulou et al., 2014), and in future work we plan to investigate this hypothesis.

Table 3
Atlas template variability prediction. By extrapolating the power law fit of the variability measured for sample sizes <11 (see Fig. 15), it is possible to estimate the variability of surface templates generated using our full dataset size.

Target age (t) [weeks]	24	26	28	30	32	34	36
Sample size (N)	22	16	22	22	22	16	18
Variability [mm]	0.114	0.159	0.164	0.237	0.271	0.396	0.403

Appendix A. Surface modelling

MR images were first segmented using an EM algorithm (Makropoulos et al., 2014), into several regions: cortical GM, WM, CSF, lateral ventricles, the deep grey nuclei, brain stem and cerebellum. EM posterior distributions were then combined to form a probability map for matter lying inside the cortical GM/WM boundary. A super-sampled binary volume B was extracted from the probability map and from this a distance map D was computed (Maurer et al., 2003); this implicitly defines the surface boundary in the image space. A negative distance indicated a voxel was inside the surface boundary, while a positive distance indicated a voxel was outside. We then extracted an explicit surface representation for the isosurface $D = 0$.

We define our surface model S as a pair $S = \{V, E\}$ where V denotes a set of vertices $\mathbf{V} = (\mathbf{x}_1, \ldots, \mathbf{x}_{|V|})$ with position vectors $\mathbf{x}_i$ and $\mathbf{E}$ denotes an adjacency matrix, where $e_{ij} = 1$ if the vertices i and j form an edge, or 0 otherwise. The cortex consists of a thin continuous sheet of GM, enclosing the WM and sub-cortical GM structures beneath. Its surface is appropriately modelled as a closed 2D manifold that is homeomorphic to a sphere, therefore S can be initialised by any convex uniformly sampled closed triangulated mesh with zero genus, which encloses a binary volume B.

The cortical surface is then estimated by iteratively deforming S towards the boundary defined implicitly by D (Algorithm 2). Each vertex $\mathbf{x}_i$ is updated in turn and its new position is given by stepping along the inward surface normal direction $\mathbf{n}_i$, with the step size γ, where

$$\gamma = \begin{cases} \gamma_{max} & D(\mathbf{x}_i) > \gamma_{max} \\ -\gamma_{max} & D(\mathbf{x}_i) < -\gamma_{max} \\ D(\mathbf{x}_i) & \text{else} \end{cases}.$$

$\mathbf{x}_i$ is only updated if a ray cast along $\gamma\mathbf{n}_i$ intersects with B, ensuring vertices only move towards the boundary. Additionally if a ray cast along $\gamma\mathbf{n}_i$ first intersects with S at $\mathbf{y}_i$ and $\gamma + \|\mathbf{x}_i - \mathbf{y}_i\| > d$, where d is the closest permissible distance between S and itself, then $\mathbf{x}_i$ is not updated to ensure the mesh does not self intersect.

After all vertices have been updated, their positions are smoothed, taking an average of the neighbouring vertex positions, ensuring each vertex moves consistently with its neighbours (Algorithm 3). To allow the surface to dynamically expand and compress without affecting the vertex density or topology of the mesh, we remesh S after each iteration, splitting long edges and collapsing short edges. Both the update and smoothing steps are then repeated until S has converged to the implicit surface $D = 0$.

Algorithm 2. Deform mesh

Data: $S = \{V, E\}$
for $i = 1$ **to** $|V|$ **do**
$\quad N = \{j \,|\, e_{ij} = 1\}$
$\quad \overline{\mathbf{x}}_i = \sum_{j \in N} \mathbf{x}_j w_j \quad (w_j \propto 1/\|\mathbf{x}_i - \mathbf{x}_j\|^2 \quad \text{and} \quad \sum_{j \in N} w_j = 1)$
for $i = 1$ **to** $|V|$ **do**
$\quad \mathbf{x}_i = \overline{\mathbf{x}}_i$

Algorithm 3. Smooth

Data: $S = \{V, E\}$
for $i = 1$ **to** $|V|$ **do**
$\quad N = \{j \,|\, e_{ij} = 1\}$
$\quad \overline{\mathbf{x}}_i = \sum_{j \in N} \mathbf{x}_j w_j \quad (w_j \propto 1/\|\mathbf{x}_i - \mathbf{x}_j\|^2 \quad \text{and} \quad \sum_{j \in N} w_j = 1)$
for $i = 1$ **to** $|V|$ **do**
$\quad \mathbf{x}_i = \overline{\mathbf{x}}_i$

Appendix B. Graph edge ratio

For a triangular surface mesh S, with V vertices, F triangles, E edges and e connections to $k - 1$ other surfaces, the ratio $E{:}e$ is given:

1. $2E = 3F$	(Number of half edges in the graph)
2. $V - E + F = 2$	(Euler's characteristic for a zero genus polyhedron)
3. $e = V(k - 1)$	(One connection for each vertex to every other surface)

$$V - E + \frac{2}{3}E \approx 0 \quad (\text{given that } E > V \gg 2)$$
$$V - \frac{1}{3}E \approx 0$$
$$E \approx 3V$$
$$E \approx \frac{3}{k-1}e$$

Appendix C. Averaging surface normals

Averaging orientations is a non-trivial problem given the periodicity of the domain. At a corresponding location on the cortical surface, it is not biologically plausible to have large variations in orientation between subjects. By assuming that the maximum deviation from the average is bounded by π, we can take the difference in orientation to be the shortest path on the unit sphere and ignore any issues due to "wrapping".

A widely used but naive approach to this problem is the "vector sum algorithm" which averages each x, y, z component of the surface normal vector, which is equivalent to minimising the squared chord length on the unit sphere associated with each observation (Olson, 2011). Minimising the arc length is a more accurate method, assuming Gaussian noise, since the length of the arc is proportional to the difference in orientation. However, this level of accuracy is unnecessary as our chosen method is robust to inaccuracies in normal direction (Kazhdan, 2005). We therefore estimate $\overline{\mathbf{N}}_i$ from the original surface normal directions associated with each point, weighted using kernel regression, treating each dimension as a separate scalar value, and normalising the results.

References

Ajayi-Obe, M., Saeed, N., Cowan, F.M., Rutherford, M.A., Edwards, A.D., 2000. Reduced development of cerebral cortex in extremely preterm infants. Lancet 356 (9236), 1162–1163. http://dx.doi.org/10.1016/S0140-6736(00)02761-6.

Alt, H., Godau, M., 1995. Computing the Fréchet distance between two polygonal curves. Int. J. Comput. Geom. Appl. 11 (9), 75–91. http://dx.doi.org/10.1007/BF02439643.

Ashburner, J., Friston, K.J., 2000. Voxel-based morphometry—the methods. NeuroImage 11 (6), 805–821. http://dx.doi.org/10.1006/nimg.2000.0582.

Brakatsoulas, S., Pfoser, D., Salas, R., Wenk, C., 2005. On map-matching vehicle tracking data. Proceedings of the 31st International Conference on Very Large Data Bases, VLDB '05. VLDB Endowment, pp. 853–864 (URL http://dl.acm.org/citation.cfm?id=1083592.1083691).

Chung, F., 1997. Spectral Graph Theory. AMS.

Chung, M.K., Dalton, K.M., Davidson, R.J., 2008. Tensor-based cortical surface morphometry via weighted spherical harmonic representation. IEEE Trans. Med. Imaging 27 (8), 1143–1151. http://dx.doi.org/10.1109/TMI.2008.918338 (Aug).

Clouchoux, C., Kudelski, D., Gholipour, A., Warfield, S.K., Viseur, S., Bouyssi-Kobar, M., Mari, J., Evans, A.C., du Plessis, A.J., Limperopoulos, C., 2012. Quantitative in vivo MRI measurement of cortical development in the fetus. Brain Struct. Funct. 217 (1), 127–139. http://dx.doi.org/10.1007/s00429-011-0325-x.

Clouchoux, C., du Plessis, A.J., Bouyssi-Kobar, M., Tworetzky, W., McElhinney, D.B., Brown, D.W., Gholipour, A., Kudelski, D., Warfield, S.K., McCarter, R.J., Robertson, R.L., Evans, A.C., Newburger, J.W., Limperopoulos, C., 2013. Delayed cortical development in fetuses with complex congenital heart disease. Cereb. Cortex 23 (12), 2932–2943. http://dx.doi.org/10.1093/cercor/bhs281.

Dubois, J., Benders, M., Borradori-Tolsa, C., Cachia, A., Lazeyras, F., Ha-Vinh Leuchter, R., Sizonenko, S.V., Warfield, S.K., Mangin, J.F., Hüppi, P.S., 2008. Primary cortical folding in the human newborn: an early marker of later functional development. Brain 131 (8), 2028–2041. http://dx.doi.org/10.1093/brain/awn137.

Field, D.A., 1988. Laplacian smoothing and Delaunay triangulations. Commun. Appl. Numer. Methods 4 (6), 709–712. http://dx.doi.org/10.1002/cnm.1630040603.

Fischl, B., Sereno, M.I., Tootell, R.B.H., Dale, A.M., 1999. High-resolution intersubject averaging and a coordinate system for the cortical surface. Hum. Brain Mapp. 8 (4), 272–284. http://dx.doi.org/10.1002/(SICI)1097-0193(1999)8:4<272::AID-HBM10>3.0.CO;2-4.

Gaglioti, P., Danelon, D., Bontempo, S., Mombrò, M., Cardaropoli, S., Todros, T., 2005. Fetal cerebral ventriculomegaly: outcome in 176 cases. Ultrasound Obstet. Gynecol. 25 (4), 372–377. http://dx.doi.org/10.1002/uog.1857.

Garel, C., Chantrel, E., Brisse, H., Elmaleh, M., Luton, D., Oury, J., Sebag, G., Hassan, M., 2001. Fetal cerebral cortex: normal gestational landmarks identified using prenatal MR imaging. Am. J. Neuroradiol. 22 (1), 184–189.

Geva, R., Eshel, R., Leitner, Y., Valevski, A.F., Harel, S., 2006. Neuropsychological outcome of children with intrauterine growth restriction: a 9-year prospective study. Pediatrics 118 (1), 91–100. http://dx.doi.org/10.1542/peds.2005-2343.

Gholipour, A., Estroff, J.A., Warfield, S.K., 2010. Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. IEEE Trans. Med. Imaging 29 (10), 1739–1758. http://dx.doi.org/10.1109/tmi.2010.2051680.

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. http://dx.doi.org/10.1016/j.neuroimage.2012.05.083.

Habas, P.A., Scott, J.A., Roosta, A., Rajagopalan, V., Kim, K., Rousseau, F., Barkovich, A.J., Glenn, O.A., Studholme, C., 2012. Early folding patterns and asymmetries of the normal human brain detected from in utero MRI. Cereb. Cortex 22 (1), 13–25. http://dx.doi.org/10.1093/cercor/bhr053.

Heckemann, R.A., Hajnal, J.V., Aljabar, P., Rueckert, D., Hammers, A., 2006. Automatic anatomical brain MRI segmentation combining label propagation and decision fusion. NeuroImage 33 (1), 115–126. http://dx.doi.org/10.1016/j.neuroimage.2006.05.061.

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. J. Neurosci. 30 (6), 2268–2276. http://dx.doi.org/10.1523/JNEUROSCI.4682-09.2010.

Jain, V., Zhang, H., 2006. Robust 3D shape correspondence in the spectral domain. Shape Modeling and Applications, 2006. SMI 2006. IEEE International Conference on, p. 19-19 http://dx.doi.org/10.1109/SMI.2006.31 (June).

Jiang, S., Xue, H., Glover, A., Rutherford, M., Rueckert, D., Hajnal, J.V., 2007. MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR): application to fetal, neonatal, and adult brain studies. IEEE Trans. Med. Imaging 26 (7), 967–980. http://dx.doi.org/10.1109/tmi.2007.895456.

Kandel, E., 2013. Principles of neural science. Principles of Neural Science Fifth edition. McGraw-Hill Education 9780071390118.

Kapellou, O., Counsell, S.J., Kennea, N., Dyet, L., Saeed, N., Stark, J., Maalouf, E., Duggan, P., Ajayi-Obe, M., Hajnal, J., Allsop, J.M., Boardman, J., Rutherford, M.A., Cowan, F., Edwards, A.D., 2006. Abnormal cortical development after premature birth shown by altered allometric scaling of brain growth. PLoS Med. 3 (8), e265. http://dx.doi.org/10.1371/journal.pmed.0030265 (08).

Kazhdan, M., Hoppe, H., 2013. Screened Poisson surface reconstruction. ACM Trans. Graph. 32 (3), 29:1–29:13. http://dx.doi.org/10.1145/2487228.2487237 (June).

Kazhdan, Michael, 2005. Reconstruction of solid models from oriented point sets. Proceedings of the Third Eurographics Symposium on Geometry Processing, SGP '05. Eurographics Association (URL http://dl.acm.org/citation.cfm?id=1281920.1281931).

Kim, J.S., Singh, V., Lee, J.K., Lerch, J., Ad-Dab'bagh, Y., MacDonald, D., Lee, J.M., Kim, S.I., Evans, A.C., 2005. Automated 3-D extraction and evaluation of the inner and outer cortical surfaces using a laplacian map and partial volume effect classification. NeuroImage 27 (1), 210–221. http://dx.doi.org/10.1016/j.neuroimage.2005.03.03.

Kim, K., Habas, P.A., Rousseau, F., Glenn, O.A., Barkovich, A.J., Studholme, C., 2010. Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. IEEE Trans. Med. Imaging 29 (1), 146–158. http://dx.doi.org/10.1109/tmi.2009.2030679.

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. Med. Image Anal. 16 (8), 1550–1564. http://dx.doi.org/10.1016/j.media.2012.07.004.

Kyriakopoulou, V., Vatansever, D., Elkommos, S., Dawson, S., McGuinness, A., Allsop, J., Molnár, Z., Hajnal, J., Rutherford, M., 2014. Cortical overgrowth in fetuses with isolated ventriculomegaly. Cereb. Cortex 24 (8), 2141–2150. http://dx.doi.org/10.1093/cercor/bht062.

Li, G., Wang, L., Shi, F., Lin, W., Shen, D., 2014. Constructing 4D infant cortical surface atlases based on dynamic developmental trajectories of the cortex. Medical Image Computing and Computer-assisted Intervention. Lecture Notes in Computer Science vol. 8675. Springer International Publishing, pp. 89–96. http://dx.doi.org/10.1007/978-3-319-10443-0_12.

Lombaert, H., Grady, L., Polimeni, J.R., Cheriet, F., 2011. Fast brain matching with spectral correspondence. Information Processing in Medical Imaging. Lecture Notes in Computer Science vol. 6801. Springer Berlin Heidelberg, pp. 660–673. http://dx.doi.org/10.1007/978-3-642-22092-0_54.

Lombaert, H., Grady, L., Polimeni, J.R., Cheriet, F., 2013a. Focusr: feature oriented correspondence using spectral regularization—a method for precise surface matching. IEEE Trans. Pattern Anal. Mach. Intell. 35 (9), 2143–2160. http://dx.doi.org/10.1109/TPAMI.2012.276 (Sept.).

Lombaert, H., Sporring, J., Siddiqi, K., 2013b. Diffeomorphic spectral matching of cortical surfaces. Information Processing in Medical Imaging (IPMI). Lecture Notes in Computer Science vol. 7917. Springer, pp. 376–389.

Lui, J.H., Hansen, D.V., Kriegstein, A.R., 2011. Development and evolution of the human neocortex. Cell 146 (1), 18–36. http://dx.doi.org/10.1016/j.cell.2011.06.030.

Lyttelton, O., Boucher, M., Robbins, S., Evans, A., 2007. An unbiased iterative group registration template for cortical surface analysis. NeuroImage 34 (4), 1535–1544. http://dx.doi.org/10.1016/j.neuroimage.2006.10.041.

Makropoulos, A., Gousias, I.S., Ledig, C., Aljabar, P., Serag, A., Hajnal, J.V., Edwards, A.D., Counsell, S.J., Rueckert, D., 2014. Automatic whole brain MRI segmentation of the developing neonatal brain. IEEE Trans. Med. Imaging 33 (9), 1818–1831. http://dx.doi.org/10.1109/TMI.2014.2322280 (Sept).

Mateus, D., Horaud, R., Knossow, D., Cuzzolin, F., Boyer, E., 2008. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, pp. 1–8 http://dx.doi.org/10.1109/CVPR.2008.4587538 (June).

Maurer Jr., C.R., Qi, Rensheng, Raghavan, V., 2003. A linear time algorithm for computing exact Euclidean distance transforms of binary images in arbitrary dimensions. IEEE Trans. Pattern Anal. Mach. Intell. 25 (2), 265–270. http://dx.doi.org/10.1109/TPAMI.2003.1177156.

Nadaraya, E., 1964. On estimating regression. Theory Probab. Appl. 9 (1), 141–142. http://dx.doi.org/10.1137/1109020.

Olson, E., 2011. On computing the average orientation of vectors and lines. Robotics and Automation (ICRA), 2011 IEEE International Conference on, pp. 3869–3874 http://dx.doi.org/10.1109/ICRA.2011.5979562 (May).

Orasanu, E., Melbourne, A., Lombaert, H., Cardoso, M.J., Johnsen, S.F., Kendall, G.D., Robertson, N.J., Marlow, N., Ourselin, S., 2014. Prefrontal cortical folding of the preterm brain: a longitudinal analysis of preterm-born neonates. MICCAI Workshop on Spatiotemporal Image Analysis for Longitudinal and Time-series Image Data (STIA).

Robinson, E.C., S.J., 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 1000 (0), 414–426. http://dx.doi.org/10.1016/j.neuroimage.2014.05.069.

Rohlfing, T., Russakoff, D.B., Maurer, C.R., 2003. Expectation maximization strategies for multi-atlas multi-label segmentation. Information Processing in Medical Imaging. Lecture Notes in Computer Science vol. 2732. Springer, Berlin Heidelberg, pp. 210–221. http://dx.doi.org/10.1007/978-3-540-45087-0_18.

Rousseau, F., Glenn, O.A., Iordanova, B., Rodriguez-Carranza, C., Vigneron, D.B., Barkovich, J.A., Studholme, C., 2006. Registration-based approach for reconstruction of high-resolution in utero fetal MR brain images. Acad. Radiol. 13 (9), 1072–1081. http://dx.doi.org/10.1016/j.acra.2006.05.003.

Rousseau, F., Kim, K., Studholme, C., Koob, M., Dietemann, J.L., 2010. On super-resolution for fetal brain MRI. Medical Image Computing and Computer-assisted Intervention: MICCAI. Lecture Notes in Computer Science: LNCS vol. 6362, pp. 355–362. http://dx.doi.org/10.1007/978-3-642-15745-5_44.

Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J., 1999. Nonrigid registration using free-form deformations: application to breast MR images. IEEE Trans. Med. Imaging 18 (8), 712–721. http://dx.doi.org/10.1109/42.796284.

Serag, A., Aljabar, P., Ball, G., Counsell, S.J., Boardman, J.P., Rutherford, M.A., Edwards, A.D., Hajnal, J.V., Rueckert, D., 2012. Construction of a consistent high-definition spatio-temporal atlas of the developing brain using adaptive kernel regression. NeuroImage 59 (3), 2255–2265. http://dx.doi.org/10.1016/j.neuroimage.2011.09.062.

Sørensen, T., 1948. A method of establishing groups of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on Danish commons. Biol. Skr. 5, 1–34.

Van Essen, D.C., 2005. A population-average, landmark- and surface-based (pals) atlas of human cerebral cortex. NeuroImage 28 (3), 635–662. http://dx.doi.org/10.1016/j.neuroimage.2005.06.058.

Van Essen, D.C., Drury, H.A., 1997. Structural and functional analyses of human cerebral cortex using a surface-based atlas. J. Neurosci. 17 (18), 7079–7102 (URL http://www.jneurosci.org/content/17/18/7079).

Wright, R., Kyriakopoulou, V., Ledig, C., Rutherford, M.A., Hajnal, J.V., Rueckert, D., Aljabar, P., 2014. Automatic quantification of normal cortical folding patterns from fetal brain MRI. NeuroImage 91, 21–32. http://dx.doi.org/10.1016/j.neuroimage.2014.01.034.

Yeo, B.T.T., Sabuncu, M.R., Vercauteren, T., Ayache, N., Fischl, B., Golland, P., 2010. Spherical demons: fast diffeomorphic landmark-free surface registration. IEEE Trans. Med. Imaging 29 (3), 650–668. http://dx.doi.org/10.1109/TMI.2009.2030797 (March).