Appendix A. Data

Human diffusion MRI data

As part of the WU-Minn Human Connectome Project (Van Essen et al. 2013) diffusion MRI data was collected for 200 subjects at 1.25 mm resolution on a 3T scanner as well as 1.05 mm resolution on a 7T scanner. For our analysis we use the pre-processed diffusion MRI data as well as the white/gray matter boundary and pial surface extracted using the HCP pipelines (Glasser et al. 2013) for 10 subjects scanned both in the 3T and 7T scanners. The HCP pipelines use a customized Freesurfer pipelines to extract these surfaces from T1-weighted and T2-weighted images (Dale, Fischl, and Sereno 1999; Fischl, Sereno, and Dale 1999). The 3T diffusion MRI data was acquired at 1.25 mm resolution covering 3 b-shells ($b = 1000$, 2000, and 3000 s/mm2) interspersed with 21 non-diffusion-weighted images ($b = 0$) (Uğurbil et al. 2013; Sotiropoulos et al. 2013). The 7T diffusion MRI data was acquired at 1.05 mm resolution covering 2 b-shells ($b = 1000$, and 2000 s/mm2) (Vu et al. 2015). During the pre-processing the diffusion MRI data were corrected for susceptibility-, motion-, and eddy-induced artifacts using FSL's topup and eddy (Andersson, Skare, and Ashburner 2003; Andersson and Sotiropoulos 2016b)

To further explore the effect of lower spatial resolutions, we analyzed additional diffusion MRI data acquired of the same subject at 3 different resolutions (1.35, 2, and 2.5 mm). The data acquisition and analysis was very similar to the HCP with 270 diffusion-weighted images obtained with $b = 1000$, 2000, and 3000 (90 volumes per shell) and 21 non-diffusion weighted images. The data was pre-processed using the same HCP pipelines.

Macaque myelin-stained data

All animal procedures were approved by the Oregon National Primate Research Center and Washington University Medical School Institutional Animal Care and Use Committees, and were conducted in full accordance with the PHS Policy on Humane Care and Use of Laboratory Animals (for details see Brambrink et al. 2012; Creeley et al. 2014).

Briefly, a six day-old (P6) infant rhesus macaque was given high-dose pentobarbital and was euthanized by transcardial perfusion-fixation with 4% paraformaldehyde solution in phosphate buffer (in keeping with National Institutes of Health guidelines) to prepare the brain for histological analysis (immunohistochemistry). The infant was naturally delivered from a dam that had been group housed or was part of the time-mated breeding program. No complications occurred during in-vivo tissue fixation.

After post-fixation for another 24 hours with 4% paraformaldehyde at 4°C, 70 µM serial coronal sections were cut on a vibratome across the entire rostro-caudal extent of the forebrain. Brain sections were used to stain myelinated axons and their trajectory in cortical white and grey matter. A monoclonal antibody to myelin basic protein was used as primary antibody (MBP, 1:100; MAB 395; EMD Millipore). Following overnight incubation with primary antibody (4°C) sections were incubated with complementary secondary antibody in blocking solution at RT and then reacted with ABC reagents (Vectastain ABC Elite Kit, Vector Labs., Burlingame, CA, USA). Vector VIP kit (Vector Laboratories, Inc., Burlingame, CA) was then used as chromogen to visualize myelinated axons. Images were photographed and digitalized with a Hamamatsu NanoZoomer 2.0-HT slide scanner.

Appendix B. Varying the interpolation scheme

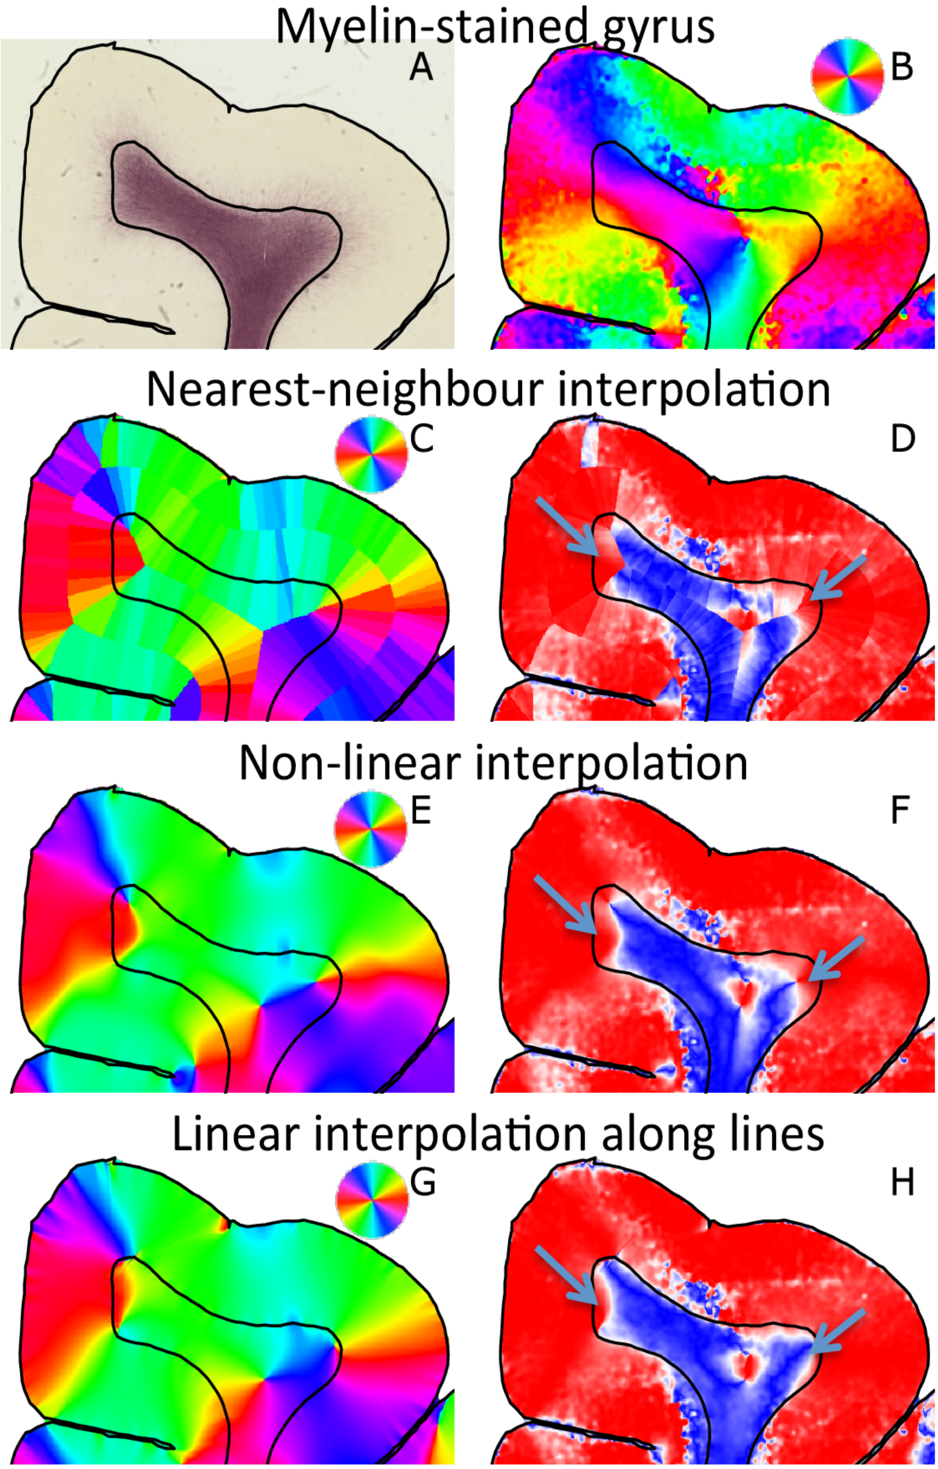

Figure B1 Myelin-stained histology section of a gyrus of a young macaque. A) raw image data; B) fibre orientations based on myelin-stained intensity; C, E, G) Radial orientation interpolated from the cortical surface in three different ways (see text); D, F, H) Alignment between observed and radial fibre orientations with radial fibres in red and tangential fibres in blue. The black lines illustrate the semi-automatically traced WM/GM boundary and pial surface. In this work we have adopted the "linear interpolation along lines" technique illustrated in panels G and H.

In this work we interpolate the surface normals and sulcal depth gradient from the WM/GM boundary and pial surface to define the gyral coordinate system in the underlying WM and cortical GM. The adopted interpolation scheme involves drawing many lines with different orientations from the point of interest to the cortical surface and linearly interpolates the normal/sulcal depth gradient along these lines. These interpolated orientations are then averaged and orthogonalized to define the gyral coordinates. Here we investigate whether similar results can be obtained using simpler interpolation schemes.

The simplest interpolation scheme is a nearest-neighbour interpolation, where the radial and tangential axes are defined with respect to the nearest surface element (McNab et al. 2013; Kleinnijenhuis et al. 2015). The radial orientations defined in this way show strong discontinuities such as seen at the gyral midplane (**Figure B1**C). Although these fibre orientations are predominantly radial in the cortex (i.e. red in **Figure B1**D) and predominantly tangential in the white matter (i.e. blue in **Figure B1**D), the observed fibre orientations are less well aligned with the gyral coordinates defined using nearest-neighbour interpolation than for the alternative interpolation schemes discussed below.

The discontinuities in the orientation of the radial axis can be reduced by not only considering the normal of the nearest surface element, but by taking an average of the normals of all surface elements. Each surface element (indexed with j) is given a weight (w_j) based on its distance from the point of interest (d_j):

$$w_j = d_j^{-\alpha}. \tag{5}$$

The average orientation is then calculated based on a PCA analysis. It is taken as the primary eigenvector of a summary matrix $A = \sum_j w_j^2\, \mathbf{n}_j \cdot \mathbf{n}_j'$ (Jones 2003), where $\mathbf{n}_j$ is the surface element normal at the surface element j. The radial axis defined using this non-linear interpolation (**Figure B1**E) is much smoother than for the nearest-neighbour interpolation, which leads to a better overall alignment of the observed fibre orientations with the radial orientation in the cortex and with the tangential orientation in the white matter (**Figure B1**F).

With both the nearest-neighbour and the non-linear interpolation introduced above the general trend of tangential fibre orientations in the white matter is broken close to the gyral crowns, where fibres run predominantly radial with respect to the gyral crown (arrows in **Figure B1**D, F). However, even in this region, fibres tend to run tangentially with respect to the sulcal walls. Hence, if we define radial orientations in this region with respect to sulcal walls rather than crowns this would increase the predominance of tangential fibre orientations seen in the white matter. This was the main argument to adopt the more complicated linear line-based interpolation scheme. A radial axis based on this scheme will tend to point radial to the sulcal walls (**Figure B1**G), which leads to more tangential orientations observed even for points close to the gyral crown (arrows in **Figure B1**H).

Although **Figure B1** suggests that this interpolation scheme might give the best alignment between the observed fibre orientations and the gyral coordinates, the real test for the different interpolation schemes is how well they align with the full 3D fibre orientations extracted from diffusion MRI data. Irrespective of the algorithm used, fibres are radial in the cortical grey matter and tangential in the underlying white matter; however, the strength of this alignment varies between the three algorithms especially in the white matter (**Figure B2**) Compared with nearest-neighbour interpolation, the adopted algorithm (linear line-based interpolation) finds more radial orientations of the primary eigenvector of the diffusion tensor in the cortex (1.3° on average) as well as more radial orientations for the tertiary eigenvector in the superficial white matter (3.0° on average). The differences with the non-linear interpolation scheme are much smaller with only an improved radiality of 0.05° on average in the cortex and of 1.4° on average in the superficial white matter. Although these improvements due to the adopted interpolation scheme may appear small, they are significant compared with the median angular offset of 7° from the tangential plane observed in the superficial white matter.

In the cortex we can also estimate the radial orientation by resampling the surfaces between the WM/GM boundary and pial surface. We extract 4 intermediate surfaces using equi-volume resampling (Waehnert et al. 2014; Kleinnijenhuis et al. 2015) and for every voxel identify the nearest surface element across all surfaces (including the WM/GM boundary and pial surface). Again a slightly poorer radiality (1.5° on average) of the primary diffusion tensor eigenvector was found (purple histogram in **Figure B2**C) compared with the proposed interpolation scheme.

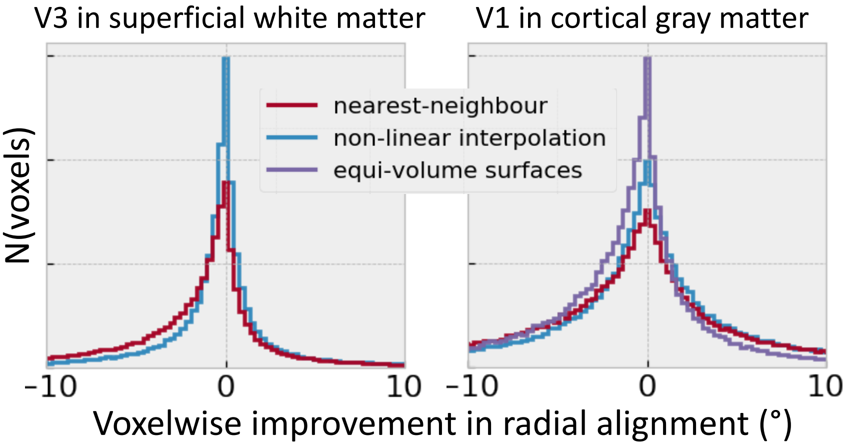

Figure B2 Histograms of the voxel-wise increase in the radiality measured for the tertiary eigenvector in the superficial white matter (left) and the primary eigenvector of the diffusion tensor in the upper cortex (right) when adopting nearest-neighbour interpolation (in red) or non-linear interpolation (in blue) rather than the proposed interpolation scheme. In the cortex we have also included a comparison with the normals from an equi-volume resampling of the surface. Positive values indicate a reduced alignment of the DTI primary eigenvector with the gyral coordinate system for these alternative interpolation schemes.

While averaging the radial orientation estimates, each estimate was assigned a weight inversely proportional to the length of the line along which it was interpolated. Varying this dependence of the weight on the line length could improve the alignment with the gyral coordinate system. To explore this we vary the power-law dependence of this relationship

$$w_i = l_i^{-\alpha}.$$

We find that either increasing or decreasing α from its adopted value of one, decreases the alignment between the radial axis and V1 in the cortex or V3 in the superficial white matter (Figure B3). Although small changes in α do not significantly alter the radiality, a larger effect is seen when assigning equal weights to all line lengths ($\alpha = 0$) or introducing a strong bias to the shortest line ($\alpha \ll 0$). The latter corresponds to the model we previously presented at ISMRM.

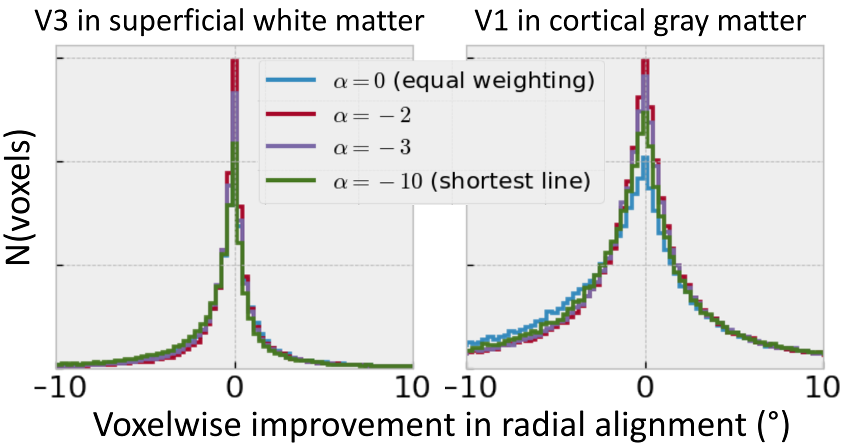

Figure B3 The voxel-wise increase in the radiality measured for the third eigenvector in the superficial white matter (left) and the first eigenvector of the diffusion tensor in the upper cortex (right) when adopting different power-law exponent (α) relating the weights and line lengths in the interpolation algorithm. Positive values indicate a reduced alignment of the DTI primary eigenvector with the gyral coordinate system for these alternative interpolation schemes.

Within the adopted linear interpolation scheme, we also need to select the number of lines N along which the normal and sulcal depth gradient are interpolated. As the number of lines decreases, the precision of the gyral coordinate estimates decreases (Figure B4). However, even at N=10 in over 80% of voxels the radial, sulcal, and gyral coordinate estimates match within 1 degree of those estimated at N=1000. For the adopted value of N=300 over 99% of voxels agree within 1 degree, which suggests that further increasing N would not change any of the results presented here.

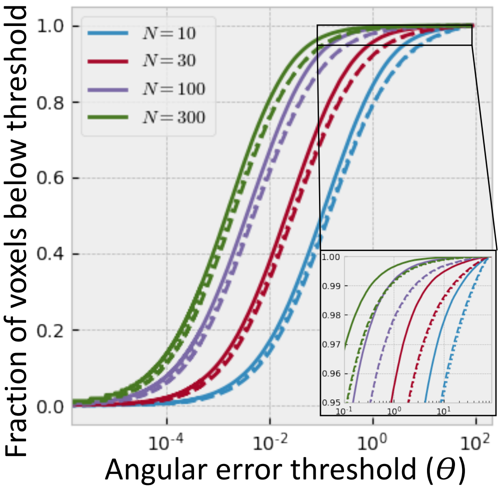

Figure B4 Comparison of the precision of estimates of the radial (solid), sulcal (dashed), and gyal (dotted) orientations for different number of lines along which to linearly interpolate (N) compared with N=1000. This plot only considers voxels in the cortex and superficial white matter.

Appendix C. Constrained spherical deconvolution

In the main text we have analysed the alignment of the gyral coordinate system with the diffusion tensor eigenvectors and the fibre orientations from FSL bedpostX. However, we can also express full orientation distribution functions (ODFs) in our gyral coordinate system.

We calculate a fibre ODF using multi-shell, multi-tissue constrained spherical deconvolution (Jeurissen et al. 2014), which has been shown to give cleaner fibre ODFs in the cortex as regular spherical deconvolution. Similar to the results presented in the main text, we find that the resulting ODF is predominantly radial in the cortex and tangential in the superficial white matter with little variability in this trend across subjects (Figure C1). In the superficial white matter this analysis confirms the predominance of sulcal orientations, although a secondary peak in the gyral orientation. This secondary peak might be caused by the decomposition of the spherical deconvolution of dispersing fibres in the tangential plane into two distinct peak rather than reflect actual crossing fibres.

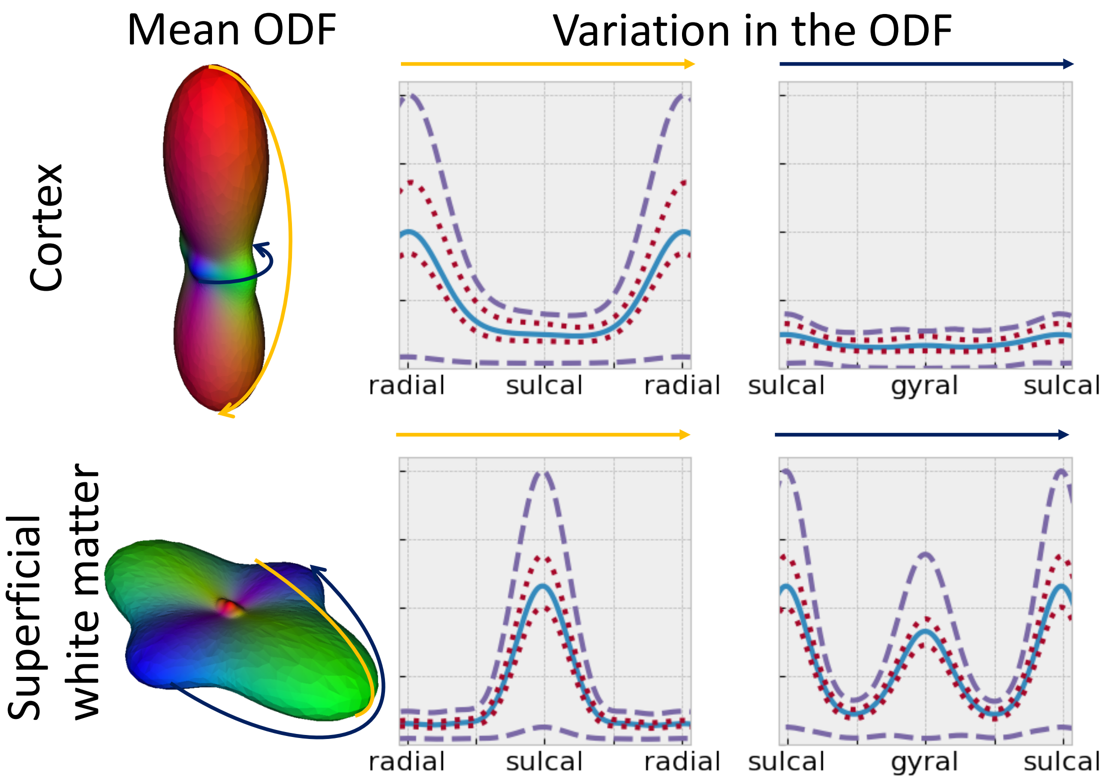

Figure C1 Left: glyphs of the mean ODF from constrained spherical deconvolution averaged over the upper cortex (> 1 mm above the WM/GM boundary) or the superficial white matter (down to 4 mm) and averaged over 29 HCP subjects. Right: The blue line illustrates the ODF averaged across subject along a line connecting the radial with the sulcal orientation (left) or the sulcal with the gyral orientation (right). The red dotted lines show the full range of the variability in this ODF across the 29 subjects. The purple dashed line shows the variability of the ODF across voxels within a single HCP subject (68% confidence interval).