

Diffusion Tensor MRI and Spherical-Deconvolution-Based Tractography on an Ultra-Low Field Portable MRI System

James Gholam¹  | Phil Schmid¹ | Joshua Ametep¹ | Alix Plumley¹ | Leandro Beltrachini²  | Francesco Padormo³ | Rui Teixeira³ | Rafael O'Halloran³ | Kaloian Petkov⁴  | Klaus Engel⁵ | Steven C. R. Williams⁶ | Sean Deoni⁷ | Mara Cercignani¹ | Derek K. Jones¹

¹Cardiff University Brain Research Imaging Center (CUBRIC), Cardiff University, Cardiff, UK | ²School of Physics and Astronomy, Cardiff University, Cardiff, Wales, UK | ³Hyperfine, Inc., Guilford, Connecticut, USA | ⁴Siemens Medical Solutions, Malvern, Pennsylvania, USA | ⁵Siemens Healthineers AG, Erlangen, Germany | ⁶Centre for Neuroimaging Sciences, King's College London, London, Camberwell, UK | ⁷Maternal, Newborn & Child Health (MNCH) and Discovery & Translational Sciences (D&T), Bill & Melinda Gates Foundation, Seattle, Washington, USA

Correspondence: James Gholam (gholamja@cardiff.ac.uk)

Received: 16 June 2025 | **Revised:** 10 October 2025 | **Accepted:** 11 January 2026

ABSTRACT

Ultra-low-field (ULF) MRI is emerging as an alternative modality to high field (HF) MRI due to its lower cost, minimal siting requirements, portability and enhanced accessibility—factors that enable large-scale deployment. Although ULF-MRI exhibits a lower signal-to-noise ratio (SNR), advanced imaging and data-driven denoising methods enabled by high-performance computing have made contrasts like diffusion-weighted imaging (DWI) feasible at ULF. This study investigates the potential and limitations of ULF tractography, using data acquired on a 0.064 T commercially available mobile point-of-care MRI scanner. The results demonstrate that most major white matter bundles can be successfully retrieved in healthy adult brains within clinically tolerable scan times. This study also examines the recovery of diffusion tensor imaging (DTI)-derived scalar maps, including fractional anisotropy and mean diffusivity. Strong correspondence is observed between scalar maps obtained with ULF-MRI and those acquired at high field strengths. Furthermore, fibre orientation distribution functions reconstructed from ULF data show good agreement with high-field references, supporting the feasibility of using ULF-MRI for reliable tractography. These findings open new opportunities to use ULF-MRI in studies of brain health, development and disease progression—particularly in populations traditionally underserved due to geographic or economic constraints. The results show that robust assessments of white matter microstructure can be achieved with ULF-MRI, effectively democratising microstructural MRI and extending advanced imaging capabilities to a broader range of research and clinical settings where resources are typically limited.

1 | Introduction

Diffusion MRI and tractography are widely used in neuroscience for developmental studies, tracking disease progression and examining tract spatial properties across various populations (Dell'Acqua et al. 2024). This broad interest in microstructural assessment and virtual dissection (Catani et al. 2002) has driven significant advances in MRI methods at high field, and in this work we seek to demonstrate the useful application of some of these techniques at ultra-low fields.

In contrast to high field (>1 T) MR systems, ULF systems (<0.1 T) are free from numerous constraints imposed by superconducting systems, including non-reliance on cryogenic cooling or high voltage power supplies, and tolerance of power supply interruptions. Alongside significant reductions in weight, B_0 fringe fields and RF heating, ULF systems often have exceptional mobility. They can be easily moved between, for example, ICU settings, remote locations or directly to emergency sites—and in many cases, they only need to be plugged into the local mains circuit or a battery supply to operate. As

This is an open access article under the terms of the [Creative Commons Attribution](#) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Human Brain Mapping* published by Wiley Periodicals LLC.

well, the reduced size, power and siting requirements and cost of ULF-systems make them attractive in research contexts, where a high-field system may represent an unacceptable cost burden to a research endeavour. The majority of contemporary diffusion-based tractography studies are conducted using data acquired on high-field MRI systems (e.g., 3 T) with moderate to high gradient amplitudes $40\text{--}80\text{ mT m}^{-1}$, enabling rapid diffusion weighted single-shot (spin) echo planar imaging (ssDW-EPI) (Tax et al. 2022). While ssDW-EPI is robust and insensitive to motion within each shot, numerous factors complicate its use at ULF. ULF DTI was first demonstrated employing non-EPI readouts (Plumley et al. 2022), and subsequently using DW-EPI (Ding 2024). The performance of fast EPI readouts are limited by low gradient strength and high B_0 inhomogeneity in ULF systems, causing significant distortions at higher resolutions and significant signal dephasing arising from low bandwidth in phase encoding directions. Long echo trains combined with slower and weaker gradients require long echo times. This produces significant T_2 -related signal loss, and while this may be mitigated by high readout bandwidth, this further degrades SNR. Additionally, at ULF, Johnson noise in inductive detector coils is the dominant noise component over body noise. This results in 3D imaging offering higher SNR at the point of detection due to the larger volume excited over 2D imaging. In contrast, the minimal SAR produced by short RF pulses at ULF allow 3D fast, multi-echo sequences to be run without exceeding SAR limits, and mitigate field inhomogeneity induced dephasing and distortion while maximising SNR.

The principal challenge hindering widespread use of ULF systems for DWI is the SNR available per unit time. As the contrast mechanism of DWI relies on signal attenuation (Jones 2010), obtaining acceptable data at the SNRs typical of ULF systems is particularly challenging. Though clinically viable ULF diffusion imaging protocols for diagnosis of cerebral ischemia and infarction (Cahn et al. 2020; Sorby-Adams et al. 2024) have recently become available, extending this to tractography is challenging as it requires high accuracy, which diagnostic DWI may not, in order to be clinically useful. Diffusion-based tractography relies upon unbiased measures with sufficient resolution to discriminate between regions containing white matter and those containing grey matter or CSF. Noise biases diffusion measurement, with Rician noise resulting in reduced apparent fractional anisotropy and mean diffusivity in DTI (Jones and Basser 2004), or, in spherical deconvolution, offsets and increased variance in volume fraction measures and relative fibre angles in mixed fibre populations (Canales-Rodríguez et al. 2015). Non-uniformity of gradient encoding and/or the static B_0 -field also complicates the unbiased interpretation of the diffusion signal.

While advances in signal detection, filtration, digitisation, reconstruction and downstream processing have enabled accelerated single-volume ULF diffusion acquisitions, significant technical obstacles remain to the implementation of a robust tractography protocol at ULF. In this work, we address some of these obstacles, and demonstrate that tractography may be performed with data acquired in feasible scan times at a magnetic field strength of 64 mT on a commercial ULF-MRI system.

2 | Methodology

2.1 | Participants

Five healthy adult volunteers were recruited for this study, with ethical approval from the Cardiff University School of Psychology ethics committee. We targeted a protocol length of 1 h (standard protocol length for research studies in our centre), additionally aiming to prevent overheating of the passively cooled gradient hardware, which is particularly stressed by the demands of diffusion-encoding pulsed gradients.

2.2 | Data Acquisition

The scanner employed was a 64 mT Swoop (hardware version 1.7, software version rc9.0 beta 1) permanent-magnet ULF-MRI system from Hyperfine Inc., equipped with a 3-axis gradient set with peak gradient amplitudes of X: 24.9 mT m^{-1} , Y: 24.4 mT m^{-1} and Z: 25.7 mT m^{-1} , at slew rates of 23, 22 and $67\text{ T m}^{-1}\text{ s}^{-1}$ respectively. Prior to imaging, the scanner was B_1 power calibrated using the manufacturer's standard pre-scan calibration, and B_0 shimmed, using linear (gradient) B_0 shimming and measured to have a B_0 homogeneity of 1100 ppm in a 16 cm diameter spherical volume. The single-yoke permanent magnet array is asymmetric in the RL-direction, producing a spatially non-uniform B_0 field.

The DWI sequence employed for this study was a 3D multi-shot diffusion weighted, non-Carr-Purcell-McGill-Bloom (CPMG), self-navigated fast-spin-echo, employing a non-Cartesian k-space trajectory, with centre-out phase encoding ordering (O'Halloran 2022). 212 shots were used for each volume, using a split echo train with 70 echoes across a TR of 800 ms, spanning 35 unique phase encodings per TR. The effective echo time was 84 ms. The field of view (FOV) was $220 \times 200 \times 180$ mm, in AP/SI/RL directions respectively. Data were reconstructed at a resolution of $3 \times 3 \times 3$ mm. The readout was fully sampled and oriented in the RL direction, whereas the phase encoding directions were oversampled by 166%. Oversampling ratios from 400% to 100% (Nyquist limit) were tested. Higher oversampling ratios were found to increase sensitivity to subject motion, manifesting as diffuse blurring and linearly increased scan time. Minimal oversampling resulted in regions of uniform signal and artificially sharp edges that did not correspond to anatomical tissue boundaries.

Diffusion sensitisation utilised a monopolar pulsed-gradient spin-echo scheme, with $b = 945\text{ s/mm}^2$ at isocentre, with diffusion encoding gradients of duration $\delta = 35\text{ ms}$ and separation $\Delta = 42\text{ ms}$. The non-diffusion weighted sequence employed $2\times$ phase oversampling to provide a high SNR for apparent diffusion coefficient (ADC) calculation, and lasted 94 s. Each diffusion-weighted volume (DWI), to collect all 212 shots, required 170 s of acquisition. Eighteen diffusion encoded volumes were acquired with isotropically distributed axes organised into three groups of six. The encoding axes in each group of six were arranged using an electrostatic repulsion algorithm (Jones and Basser 2004), enabling early scan termination if necessary while still permitting N estimation of the full diffusion tensor.

Within the same protocol, a T_2 -weighted split-echo fast-spin echo sequence was collected at a resolution of $2 \times 2 \times 2$ mm across the same FOV to provide complementary structural information to diffusion measures. The TE was 180ms, the TR was 1600ms, readout bandwidth 120kHz and echo train length 64. The data were reconstructed, denoised, debiased and distortion corrected, using the manufacturer methods. The sequence duration was 9 min 15s, and these data (which placed less demands on the gradient hardware than DWI) were collected amongst the DW volumes to provide an opportunity for the gradient hardware to cool.

Between every set of 2 DWIs, the scanner centre frequency f_0 was recentered following a short calibration scan to account for B_0 drift induced by gradient heating. Images were reconstructed using the vendor-provided reconstruction. Electromagnetic interference is removed from raw k-space data using a Kalman-filter-based algorithm and measurements acquired concurrently from 8 noise detectors. The filtered k-space data was transformed to image space using a Plug-and-Play FISTA (Sinha and Chaudhury 2024) reconstruction algorithm with a DL-based denoiser. Its data consistency step incorporated estimated 3D navigator maps into the k-space encoding operator to compensate motion.

Due to eddy-current based artifacts, the even and odd echoes of the echo train were reconstructed separately. After reconstruction, the 2 echo parities and 8 signal channels were combined in the complex domain Adaptive Reconstruction of MRI array data (Walsh et al. 2000; Griswold et al. 2002). The combined magnitude image was then corrected for geometric distortion due to gradient non-linearity and B_0 inhomogeneity, and also for the bias field based on a calibration measurement made with a flood fill phantom.

The models were trained using images from both high field and low-field MRI acquisition data across a range of image contrasts and indications. The training images were passed through a synthetic image generation pipeline that included a range of spatial, intensity and geometrical transformations to increase the variability of the training set. Finally, retrospective k-space data was generated from the augmented training image dataset by applying the forward operator to each example, simulating the effects of for example multicoil acquisition, random noise and phase. Training was performed by minimizing the L1 loss between the reconstructed image and the original target image.

The built-in bias correction was disabled to facilitate estimation of the bias field using the $b=0$ s/mm² image. The diffusion images were denoised by the vendor deep-learning algorithm separately to reconstruction.

As a result of the highly inhomogeneous B_0 field, in the scanner employed, a large permanent magnetic field gradient was present in FOV. Modelling of this gradient predicted a significant contribution to the diffusion encoding, which if characterised could be retrospectively corrected for. A prospective fieldmap was obtained prior to in vivo scanning using a gradient recalled echo image of a flood fill phantom, measured at 2 different echo times to characterise the major B_0 field variations (Figure 1).

2.3 | Data Analysis

2.3.1 | B_1 -Bias Field Correction

Reconstructed DWIs were first corrected for B_1 bias by estimation of the field on the $b=0$ s/mm² image using the N4 bias field correction algorithm (Tustison et al. 2010), and the inverse of this field applied to the $b=0$ s/mm² and all DWIs in the series (Tax et al. 2022).

2.3.2 | Correction for Background Field Gradient (B_2 Inhomogeneity)

Images were retrospectively corrected to account for the large permanent magnetic field gradient, using the $b=0$ s/mm² image as a reference to separate encoding from other contrast effects (see Appendix A). In the scanner employed in the current work, at positions 8 cm from the isocentre, this gradient may be up to 1.4 mT m^{-1} , that is up to 7% of the specified diffusion encoding gradient. This can lead to errors of 16.1% in the ADC. The contribution of this field inhomogeneity may be modelled as an additive magnetic field gradient producing a spatially dependant linear scaling to the prescribed b value by a factor $a(\mathbf{r})$. The diffusion-weighted signal S' from a voxel with apparent diffusion coefficient (ADC) D , with the contribution from the spurious field may be modelled as

$$S'(\mathbf{r}) = S_0(\mathbf{r}) \exp(-a(\mathbf{r})bD(\mathbf{r})) \quad (1)$$

with a correction depending on a and the signal from the $b=0$ s/mm² image S_0 . The correction may be represented as a simple multiplication:

$$S(\mathbf{r}) = \exp\left(\log\left(\frac{S'(\mathbf{r})}{S_0(\mathbf{r})}\right)\left(1 - \frac{1}{a(\mathbf{r})}\right)\right), \quad (2)$$

where $S(\mathbf{r})$ is the corrected signal at a given position.

2.3.3 | Registration and Super-Resolution

The diffusion-weighted volumes were affinely registered to the non-diffusion-weighted volume using the hierarchical ANTs implementation using a mutual information cost function, and the rotation component of the affine transform applied to the respective volumes b -vectors (Leemans and Jones, 2009). The structural T_2 weighted scan was registered to the $b=0$ s/mm² image, then super-resolved (SR) to a T_1 weighted contrast using SynthSR (Iglesias et al. 2021, 2023; Iglesias et al. 2022) to $1 \times 1 \times 1$ mm resolution. This structural image used as the basis for masking, partial volume calculations and in anatomically constrained tractography (Smith et al. 2012).

2.3.4 | Tissue Segmentation

The SR T_1 -weighted image was segmented using Freesurfer's recon-all pipeline (Collins et al. 1994; Dale et al. 1999; Fischl et al. 1999, 2001, 2002, 2004) to give masks of all major structures in the brain in the same space as the diffusion data. The

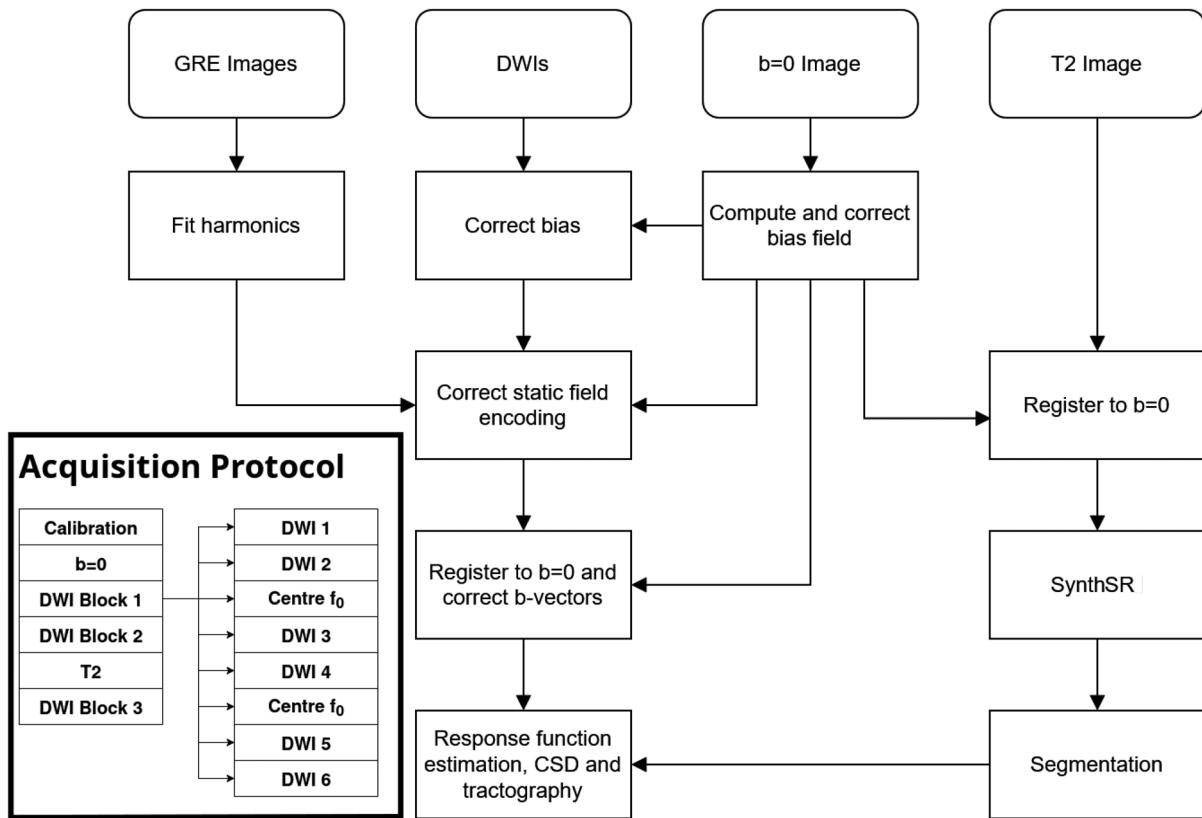


FIGURE 1 | Top right, Schematic description of the processing pipeline employed to generate corrected and combined anatomical and DW images used for subsequent tractography (indented bottom left). Acquisition scheme for imaging showing pre-scan calibration, DWI acquired in three blocks, with the subdivisions within each block shown (right of indent).

high-resolution information afforded by this approach is desirable for the creation of grey- and white-matter masks that allow inference of the partial volumes present in the lower resolution diffusion data. A 5 tissue-type (5tt) segmentation was constructed from this, as was a grey-matter white-matter interface (GMWMI) mask. Manual voxel selection in the corpus callosum midline was used to estimate single-fibre response functions with spherical harmonic order 4 (to match the number of equations solved by constrained spherical deconvolution (CSD) to the number of DWIs) using mrtrix3 (Tournier et al. 2019).

2.3.5 | DTI, Fibre Orientation and Tractography

Diffusion tensors were fitted to the observed data using the RESTORE method (Chang et al. 2005; Chung et al. 2006; Yendiki et al. 2014) implemented in DIPY Version 1.11 (Garyfallidis et al. 2014). Noise levels were initially estimated with the DIPY “estimate_sigma” method (Koay and Basser 2006; Park et al. 2009), then manually adjusted to maximise visual consistency.

The response functions permitted estimation of fibre-orientation distributions for each tissue type using the MSMT-CSD (Jeurissen et al. 2014) method given in mrtrix3. This was then used alongside the 5tt masks and the GMWMI masks to perform anatomically constrained tractography using the mrtrix3 implementation of the iFod2 algorithm. Final tractograms were computed using three methods:

1. Manual definition of inclusion and exclusion ROIs for major WM bundles, using the super-resolved T_1 to estimate anatomically informed regions of interest (Catani et al. 2002; Conturo et al. 1999). These ROIs were used to filter a 2 million streamline wholebrain iFod2-ACT tractogram, only including fibres within inclusion ROIs, and trimming fibres within exclusion ROIs.
2. Filtering of a 10-million streamline iFod2-ACT tracked brain using TractSeg ROIs. Bundle ending segmentations (i.e., spatial estimates of the cortical terminations of WM bundles) were generated using the tool TractSeg (Wasserthal et al. 2018a), using previously computed fODF peaks in WM. Only tracts that started and ended in bundle endings, and remained within bundle ROIs were retained, using the corresponding ROIs generated using TractSeg.
3. Tract-orientation mapping (TOM) (Wasserthal et al. 2018b) using TractSeg derived ROIs and fODF peaks in WM.

2.3.6 | Visualisation

The ultra-low-field MRI data was visualised using Cinematic Rendering (Comaniciu et al. 2016; Dappa et al. 2016), a Monte-Carlo path tracing engine developed by Siemens Healthineers, which integrates various data sources from medical imaging to generate photorealistic images and animations. Fibre data was processed with 3D Slicer (Fedorov et al. 2012) and MRtrix3 (Tournier et al. 2019), transforming MRtrix.tck files

into VTK (Schroeder et al. 2006) polygonal data. Fibre probability density maps were generated using MRtrix3's tckmap command, producing directionally encoded colour space fibre data and fibre density volume data. Polygonal fibre data was integrated with T2 volume data using a unified path tracing approach, constructing a bounding volume hierarchy for fast intersection of rays with fibre representations. The Cinematic Rendering engine registers T2-weighted MRI sequences with volumetric fibre representations, applying a brain mask signed distance field to reveal anatomical structures. Users can adjust fibre data density and colour through interactive transfer functions. During path tracing, samples from all volumes and meshes are composited in the shader, using high-dynamic range lighting and tone-mapping for final output. The polygonal fibre method generates many primitives, limiting visualisation to specific bundles, while the volumetric method allows visualisation of the entire fibre set, offering sharper representation and comprehensive coverage. Both methods can be combined for comprehensive coverage while highlighting selected fibre bundles with a precise polygonal representation.

2.3.7 | High-Field Validation Dataset

High field measurements were also collected to serve as a high SNR reference standard against ULF measurements. Data were collected on the same participants as for ULF at 3T on a Siemens Connectom MR system with 300mTm^{-1} gradients using a multiband DW-echo-planar imaging sequence, following a high-angular-resolution diffusion imaging (HARDI) protocol with 253 directions at b-values of 0, 200, 500, 1200, 2400, 4000 and 6000s mm^{-2} with shells containing 13, 20, 20, 30, 61, 61 and 61 isotropically distributed directions respectively. Imaging resolution was $2 \times 2 \times 2\text{ mm}$, TR 3000ms, TE 59ms, FOV $220 \times 220 \times 132\text{mm}^3$. Diffusion encoding used gradients of duration $\delta = 7\text{ms}$ and separation $\Delta = 24\text{ms}$. The total readout duration was 29ms. Data were corrected for drift using an in house method (Vos et al. 2017), and similarly for gradient non-uniformity induced distortions. Data were further corrected for susceptibility and eddy current induced distortions using a reverse-phase encoding method (Andersson et al. 2003, 2016; Andersson and Sotiroopoulos 2016; Smith et al. 2004), and denoised with MP-PCA (Veraart et al. 2016). Gibbs deringing using a subvoxel shift method was additionally employed (Kellner et al. 2016). A T_1 weighted MP-RAGE structural scan was acquired and processed with Freesurfer's recon-all and mrtrix3 to produce a 5tt segmentation and corresponding grey-matter white-matter interface masks, and these were affinely coregistered to the diffusion data.

The diffusion data were used to compute fODFs, mirroring the methods employed at ULF, though with maximum spherical harmonic $l=8$, and the Dhollander algorithm (Dhollander et al. 2016) was used to estimate the tissue response functions. Lastly, fODFs were also calculated using MSMT-CSD, restricting the b values to those above 2400s/mm^2 to benefit from the improved resolution of orientation offered at these weightings. As well, 2 $b=0\text{s/mm}^2$ and 30 $b=1200\text{s/mm}^2$ volumes were extracted from the data and fitted to spherical harmonic $l_{\text{max}}=4$, identically to the ULF data. This high SNR but low

direction count dataset was used to inform on the impact of SNR on fODF measures versus the ULF data. The high direction count and high b-value dataset was used as a “silver standard” measurement.

3 | Results

3.1 | Diffusion and Fibre Orientation Measures

Diffusion measures at ULF visually corresponded with HF measures taken with small numbers of directions and a single shell. Diffusion encoded colour maps (Figure 2A,B) showed good agreement in major tracts; however, anisotropic tissue further from the mid-sagittal axis plane (and hence isocentre) was not fully recovered, likely due to the high noise floor suppressing apparent fractional anisotropy. This was more visible in ULF FA maps directly (Figure 2C,D), where subcortical fibres were generally not observed, though were highly visible at HF. Imperfection and inconsistency in the shape of some major fibre bundles was observed, for example in the major forceps, where one side was measured to be more anisotropic than the other. Mean diffusivity maps (Figure 2E,F) showed generally good correspondence. Image quality and subsequent DTI measures and tracking quality were not found to vary substantially between individuals, and the results shown herein are a representative result of a compliant subject. All participants were compliant and tolerated scanning of this duration well, and no subjects displayed any badly motion-corrupted volumes, with the navigated multi-shot DWI employed offering tolerance to isolated incidental movements.

Fibre orientation measures were also broadly comparable between HF and ULF (Figure 3), though limitations in effective resolution were marked, with large fODF lobes spuriously attributed to grey matter and CSF. Reducing the number of directions and shells available to the HF measurement (Figure 3B) demonstrated that the quality of fODF estimation is highly dependent on large numbers of diffusion directions and multiple shells. Despite these limitations, CSD correctly estimated the orientation of some subcortical fibres, even though DTI metrics failed to assign anisotropy to these regions.

3.2 | Tractography

Tractography derived from orientation maps using manual ROI selection showed good coherence in the core of the fasciculus (where anisotropy is high), but rapid dispersion when moving into more cortical regions (as in Figure 4D). Gradual curvature within tracts such as the superior longitudinal fasciculus, arcuate fasciculus and corpus callosum (Figure 4B,F) was correctly recovered, suggesting that diffusion data has sufficient SNR to produce stable orientation estimates.

A significant fraction of major WM tracts were reconstructed, including the arcuate, inferior longitudinal, inferior fronto-occipital, uncinate and superior longitudinal fasciculi. As well, the cingulum, the upper portion of the fornix and the pontine and corticospinal tracts were well reconstructed. The major interhemispherical corpus callosum was

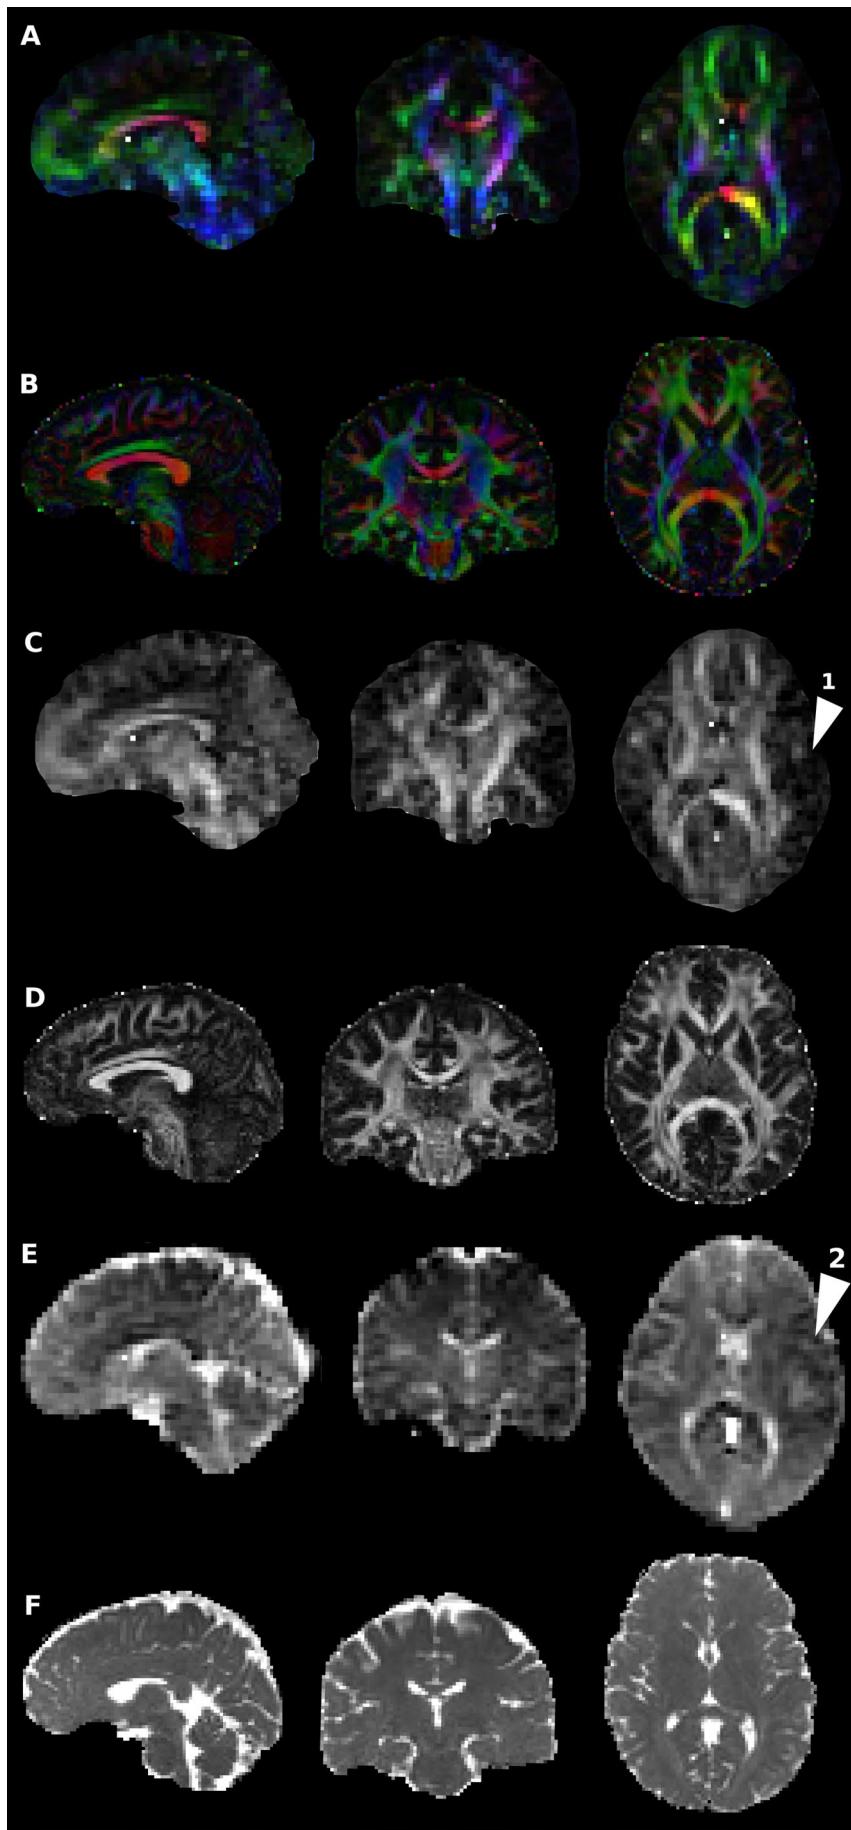


FIGURE 2 | Legend on next page.

FIGURE 2 | Comparison of quantitative diffusion between high- (B, D and F) and ultra-low-field (A, C and E) measures obtained in the same subject in sagittal, coronal and axial planes. High field measures used 30 directions at $b = 1200 \text{ s/mm}^2$. (A, B) Diffusion encoded colour maps, weighted by FA. Distinct asymmetry is observed between left and right hemispheres in ULF data (arrow 2). This may arise from different spatial noise dependence, or from uncorrected diffusion encoding nonuniformity. (C, D) Fractional anisotropy—FA is noticeably lower in ULF measurements, again likely due to the elevated noise floor. Transverse WM is distinctly darker in FA measurements at ULF than in HF (arrow 1), possibly due to the broad PSF and partial volume effects causing thinner bundles to appear more diffuse. (E, F) Mean diffusivity—shading is observed in ULF measurements, but otherwise there is good correspondence between HF and ULF measures.

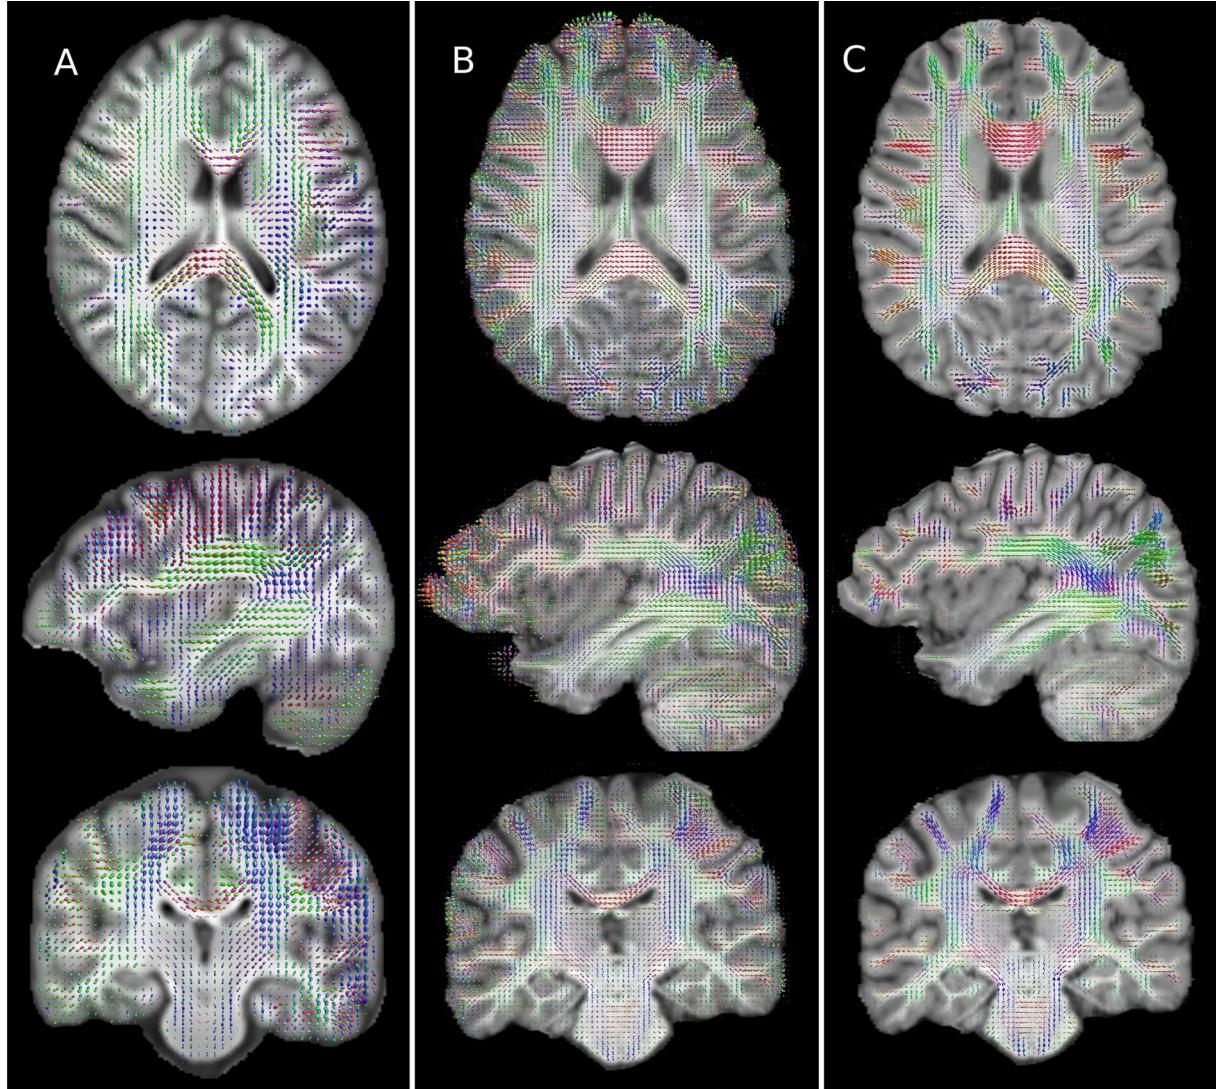


FIGURE 3 | Comparison of fibre orientation distribution plots computed in the same representative subject for the ULF (A), HF with 30 directions and one shell (B), and HF with 253 diffusion volumes across 6 shells (C). Substantial blurring is apparent in A, with margins of WM structures being less well defined and adjacent structures averaging together. Conspicuously this leads to large assigned fibre populations in GM, particularly in highly convoluted regions. Comparing A with C shows there is generally good correspondence between ULF and the gold standard measurement. Comparing B and C shows the significant impact of a reduced direction count on an otherwise high SNR measurement, with B showing smaller bundles having erroneously large lobes, and grey matter with significant anisotropy that is not corroborated by C.

generally well recovered but showed pronounced morphological differences to expected anatomy in the proximity of the ventricles, where motion artifacts are expected to be more severe.

Application of deep-learning priors through TractSeg resulted in much smoother fibre trajectory reconstruction that conformed closer to prior expectations of white matter anatomy (as in

Figure 5), and qualitatively resembled results obtained using the semi-automated approach using the same ROIs. Tracts obtained using TOM were highly coherent, but the major trunks of bundles still followed approximately the same courses as for CSD-ACT (see Figure 6). Tracts such as the SLF were reconstructed well by TOM (see Figure 7), but semi-automated approaches did not perform as well in this large tract with numerous complex features.

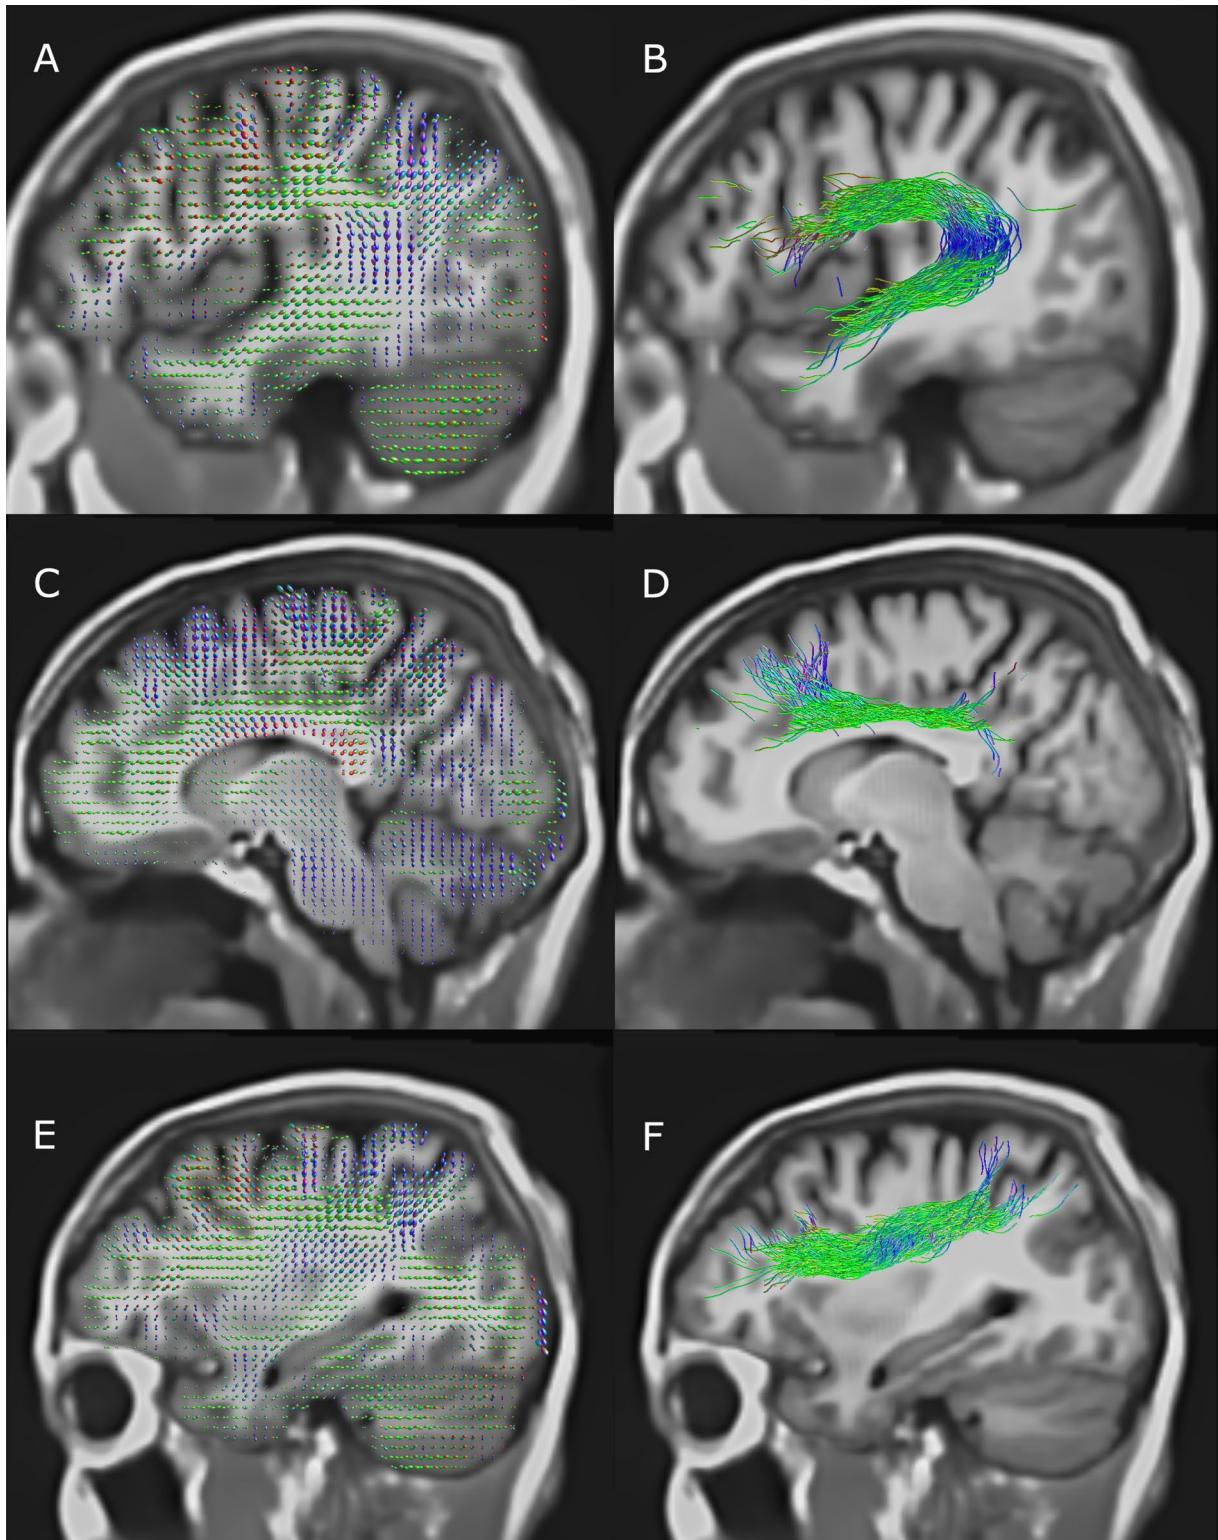


FIGURE 4 | Fibre ODFs and major association tracts retrieved at ULF using manual ROI selection; (A, B) arcuate fasciculus; (C, D) upper portion of cingulum; (E, F) SLF I and II. Tracts shown are not cropped to the displayed slice.

TractSeg generates vector-valued data at each voxel, for each tract considered. This is generated using DL segmentation of fibre bundles alongside the voxel-wise vectors describing the orientation of the three largest spherical harmonic peaks in an input fODF. Consequently, TOM may be viewed as a tract-specific synthesis of the three largest fibre populations in each

voxel based on learned priors. For large tracts with dominant single-fibre regions, it is expected that the largest SH peak corresponds to the tract orientation. Figure 9 shows the orientation of the CSD fODF peaks alongside the orientations predicted by TOM, as well as the angular difference between the two. This is shown for two large tracts with significant single-fibre content,

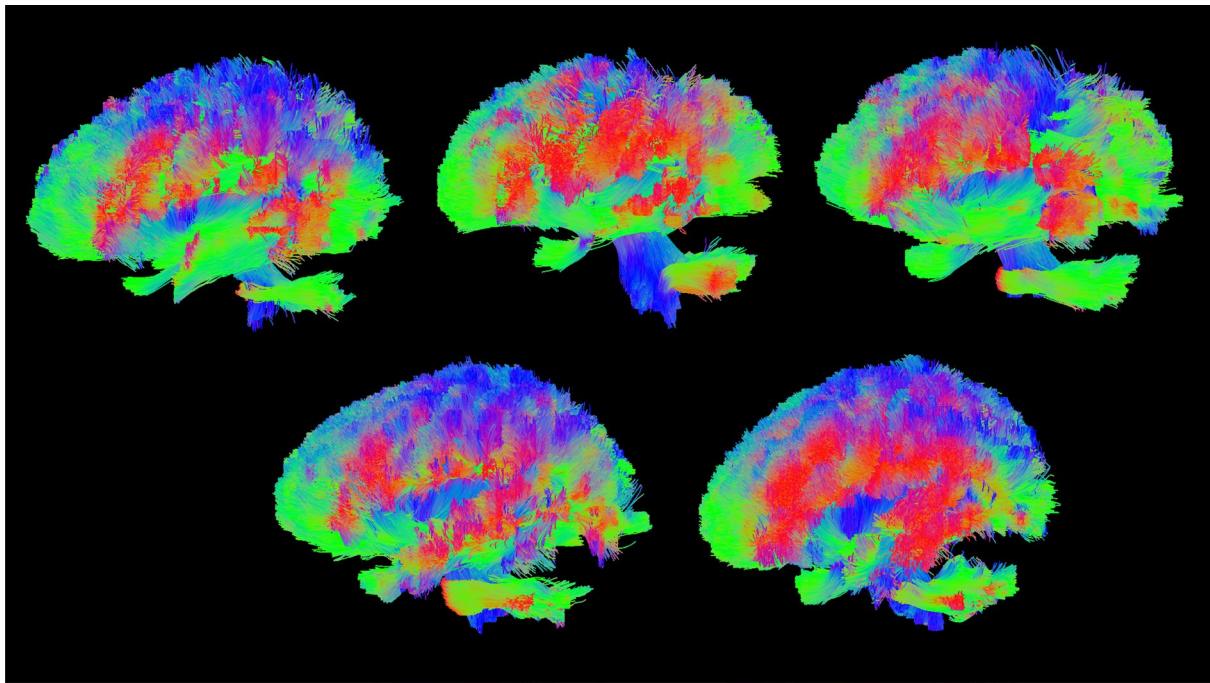


FIGURE 5 | Whole-brain automated tracking conducted with TractSeg for the five subjects imaged during this study. Major tracts were consistently recovered for all subjects, however some smaller tracts were not.

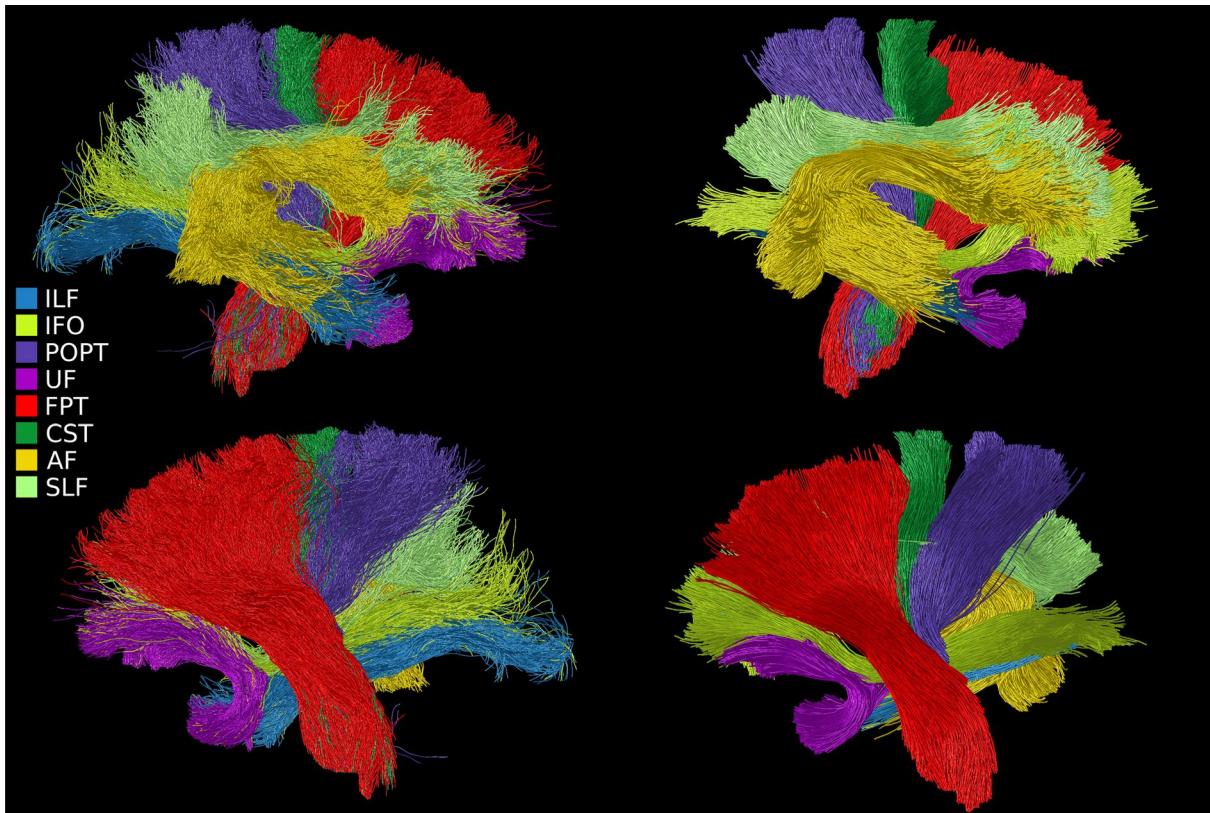


FIGURE 6 | Selection of major WM tracts in a single hemisphere of the brain retrieved using automated ROI filtering of the wholebrain tractogram generated with iFoD2 and ACT (left column) and TOM tracking with TractSeg (right column). Visible tracts are the inferior longitudinal fasciculus (ILF), inferior occipito-frontal fasciculus (IFO), parieto-occipital pontine (POPT), uncinate fasciculus (UF), fronto-pontine tract (FPT), corticospinal tract (CST), arcuate fasciculus (AF) and superior longitudinal fasciculus (SLF). TOM and automated filtering show close correspondence, though TOM measures are smoother, and truncated upon reaching the cortex.

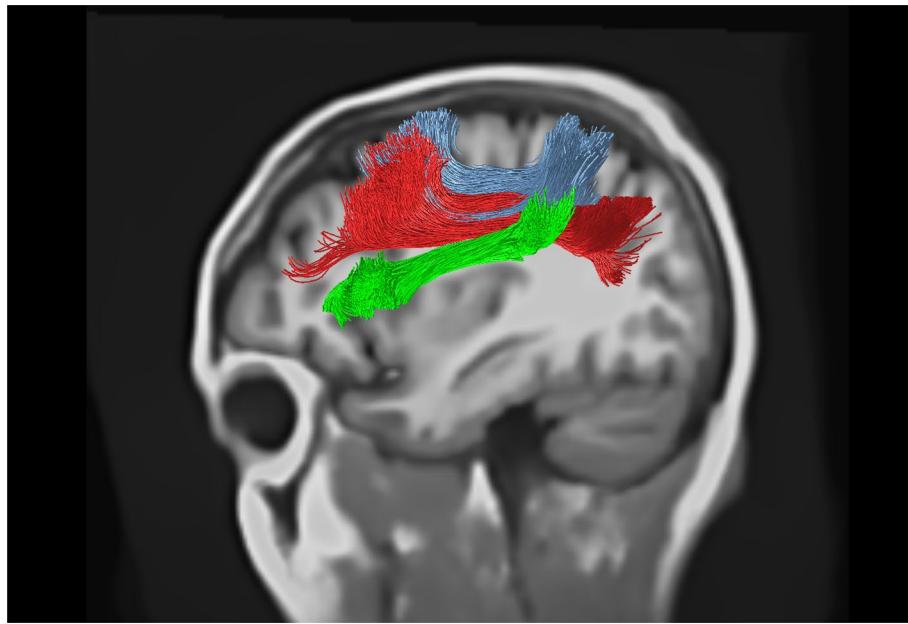


FIGURE 7 | TractSeg-derived tracking of the superior longitudinal fasciculus, reconstructing the SLF I (blue), II (red) and III (green) bundles.

the corticospinal tract and the arcuate fasciculus. Deviations are below 20 degrees in the majority of the trunk of both tracts, with errors increasing in tissues closer to the cortex.

Tracking was broadly consistent across scans, with similar quality for all 5 subjects (Figure 5). TractSeg was always able to segment and assign tracking to the corticospinal tract, the arcuate, longitudinal, medial lateral, inferior occipito-frontal and uncinate fasciculi, both cingula, the anterior thalamic radiation and the corpus callosum. As well, most striatal and thalamic tracts were retrieved, along with the cerebral peduncles, though not all segments were recovered. The fornix and commissure anterior were never retrieved by TractSeg.

3.3 | Artifacts

Significant artifacts were observed in all ULF scans conducted, and were of varying consequence in subsequent post-processing. Major artifacts and their remediation (if available) were:

- Partial voluming arose from the large voxel size employed in this study—this was controlled through the application of the multi-tissue CSD technique, allowing large voxels to remain informative of fibre behaviour.
- Diffusion encoding nonuniformity arising from the large inhomogeneous B_0 field gradient relative to the encoding gradients. This was addressed with the magnitude correction approach outlined in Methods and Appendix A.
- Broad point spread functions from T_2 attenuation during long readouts—these were addressed solely by measures to shorten the readout duration.
- Smoothly spatially varying intensity modulation, which differed between volumes, possibly due to eddy currents or motion. These were not addressed in the present study.

Partial voluming in the large voxels used is inevitable, an issue exaggerated by movement during the long acquisition. The effect of this is apparent in the reduced fractional anisotropy in narrow fibres measured at ULF vs. HF, as well as in defects visible in for example the splenium where the narrow tract coincides with the adjacent ventricle (see Figure 2). Regions with high CSF content showed apparent fibre content inconsistent with high field reference scans and known anatomy (e.g., the midsagittal surface as in Figure 8A).

The poor spatial localisation was exacerbated by the broad PSF displayed by FSE sequences with high ETL and low bandwidth, and resulted in FA values near to zero in fibres distant from the medial plane, as well as ambiguous estimates of the principal orientation of the diffusion tensor (as in Figure 2A). As well, fODFs in thinner tracts such as the cingulum were visibly diverted by signal contamination from adjacent voxels (Figure 3).

Major artifacts were observed in ODF maps in the form of spatially nonuniform responses, and are immediately visible as an asymmetric response between hemispheres (see Figure 8A). This varies spatially over relatively large length scales (15–40 mm). This artifact also occurs in DTI-based diffusion-encoded colour (Figure 2C) and mean diffusivity maps (Figures 2E and 8B), suggesting that nonuniformity in response is not related to multi-tissue CSD. This artifact resulted in a pronounced drop in tract density within the affected region (Figure 8, lower left and right).

4 | Discussion

This work demonstrates that tractography is achievable even on portable, ULF point-of-care systems. Furthermore, it establishes the viability of 3D diffusion-weighted fast spin echo (DW-FSE) as a tractography sequence at ULF, overcoming

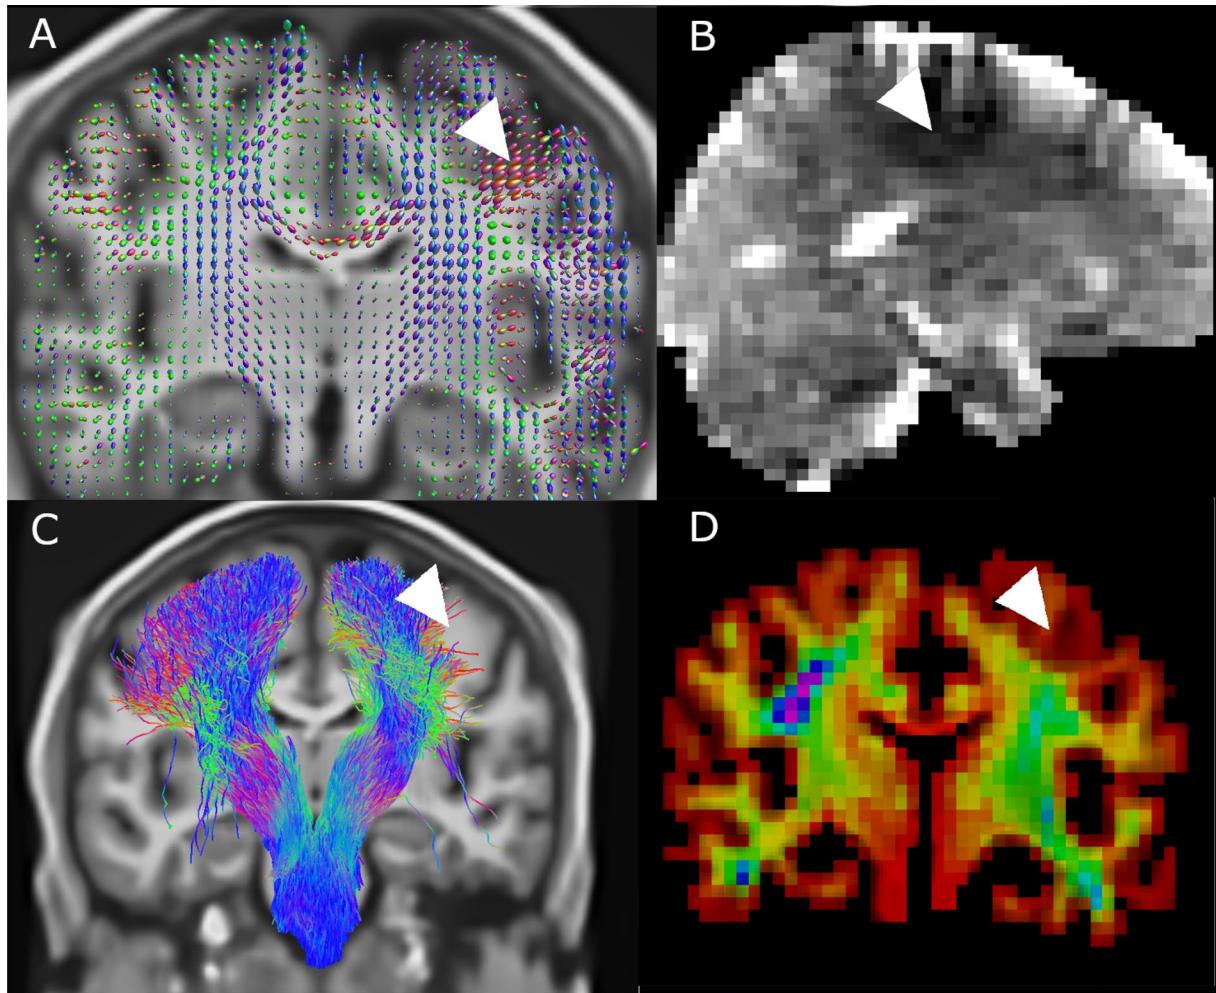


FIGURE 8 | Localised artifact (white arrow) in fODF maps (A) and corresponding hypointensity visible in mean diffusivity images (B). The hypointensity visible in MD maps is diffuse but well localised, suggesting this defect may be readily characterised and corrected. The artifact produced pronounced asymmetry in iFOD/ACT tracking of the CST (C), and corresponding wholebrain track density maps (D) corroborate the defect.

significant challenges such as low signal-to-noise ratio (SNR), poor magnet homogeneity and long, motion-sensitive readouts. The utility of ULF-derived tractography is likely in providing lower-cost, low SAR research imaging at scale, for instance to provide tract specific quantitative measures suitable for measurement at ULF such as magnetisation transfer (Balaji et al. 2025) and quantitative T_2 and T_1 (Lena 2025) metrics, as well as diffusion measures that are available by nature of the experiment.

While we report the best results to date, errors in computed mean diffusivity, fractional anisotropy and ODF scale and shape parameters highlight that artifacts as well as noise may confound downstream metrics. More artifact-immune sequences employing, for example, more sophisticated navigation or motion correction may reduce these issues to being restricted to noise alone. While the results demonstrated are encouraging, they also highlight the scale of the challenge facing ULF-based tractography. Many tools designed for high-field diffusion pre- and post-processing are employed in this work and found to translate well to ULF (such as registration, bias correction, automated segmentation and multi-tissue CSD), whereas other tools (such as nonlinear registration schemes and Gibbs deringing) were found to work inconsistently.

Several techniques employed in diffusion registration were tested in the development of this pipeline. Affine registration was found to be effective and robust, with common hierarchical approaches functioning as intended. The low resolution of the base DWIs masked poor registration to a significant degree, and indications of thermal drift producing nonidentical deformations were observed. The low consistency between DWIs from for example variation in the presentation of motion or eddy current artifacts likely exacerbated the other issues with registration. Registration of structural images to FA template images was found to produce non-overlapping skull outlines. Nonlinear registration schemes resulted in non-smooth normalisations, particularly in the skull and brain boundaries. Highly regularised nonlinear registration to enforce the constraint of long-range spatially smooth distortions may improve the robustness of these methods.

Signal hyperintensities in the DWI data led to underestimation of mean diffusivity and issues with fODF estimates. The broad spatial extent of these effects points to several potential sources. One candidate is magnetic field nonuniformity—either in the gradient fields, B_1 , or B_0 —which typically varies smoothly across the image. Another possibility is the presence of uncorrected gradient moments resulting from eddy currents. These

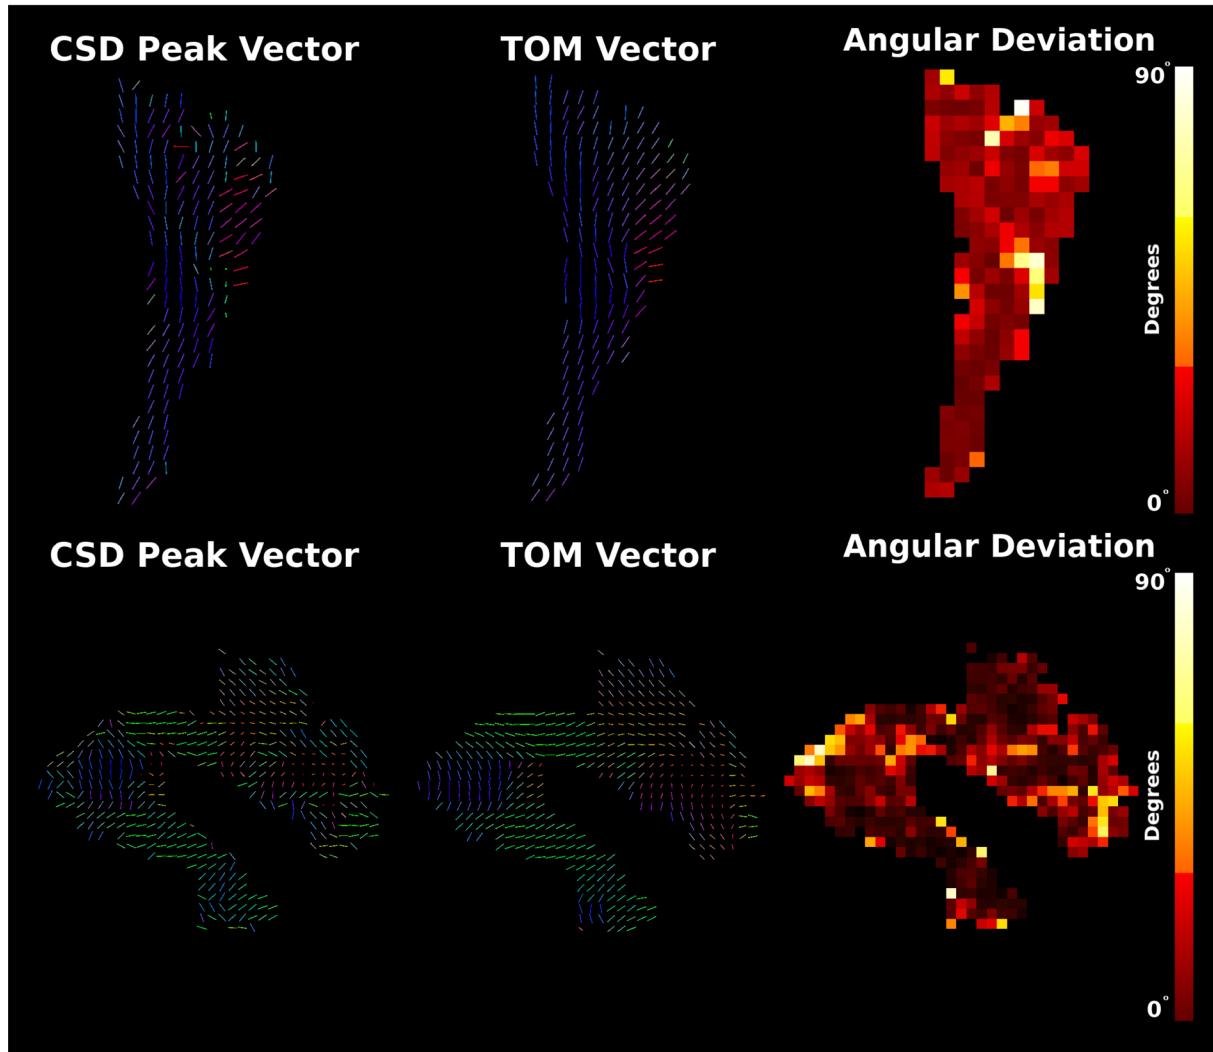


FIGURE 9 | Visualisations of the angular perturbation produced by tract orientation mapping for two major tracts, the corticospinal tract (top row) and the arcuate fasciculus (bottom row). The left column shows the voxel-wise map of the direction of the largest spherical harmonic peak (which is used by TractSeg as a basis for TOM alongside higher-order peaks), while the middle column shows the revised, tract-specific orientation estimate computed by TOM. The rightmost column shows the angular difference between the two pseudovectors.

can induce signal loss through residual gradient spoiling, with their impact expected to vary depending on the relative orientation of the diffusion-encoding direction and the readout direction. Notably, the artifact produces directional biases in the diffusion signal that are most pronounced in the lateral directions—coincidentally, also the readout direction—suggesting a systematic interaction between the acquisition scheme and underlying hardware imperfections.

Constrained spherical deconvolution proved highly effective as a model of fibre orientation vs. tensor approaches (Plumley et al. 2022), even when employing relatively sophisticated tensor fitting schemes such as RESTORE (Chang et al. 2005, 2012), and permitted the occurrence of crossing and dispersion fibres, which contributed to the quality of the tractography possible.

Despite the apparent effectiveness of CSD, the use of low b -values significantly constrains the effective resolution of fibre orientation, particularly in cases with crossing fibres with low

incidence angles (Tournier et al. 2008, 2013). Comparing CSD as a measure of orientation vs. DTI, CSD generally outperforms DTI even at a b value of $1000\text{s}\text{mm}^{-2}$ (Calamuneri et al. 2018), and is able to resolve crossing fibres at an 45° angle even at this b value for a 80-direction scheme at an SNR of 35 (Tournier et al. 2008). Our substantially lower direction count and SNR likely significantly impinges on both the accuracy of fibre orientation estimates, and of the resolution of multiple fibre populations in voxels with complex content. Despite this, prior work demonstrates successful resolution of crossing fibres with high incidence angles even in this challenging condition (Gholam et al. 2024). This is particularly likely to impact complex and smoothly diverging fibre bundles such as those in the association pathways depicted in Figure 7. In these, separation of white matter pathways is likely driven more by anatomical priors than directly by CSD.

Compatibility of the data with automated tractography approaches would be desirable, as it standardises an otherwise subjective workflow, as well as greatly speeding up an

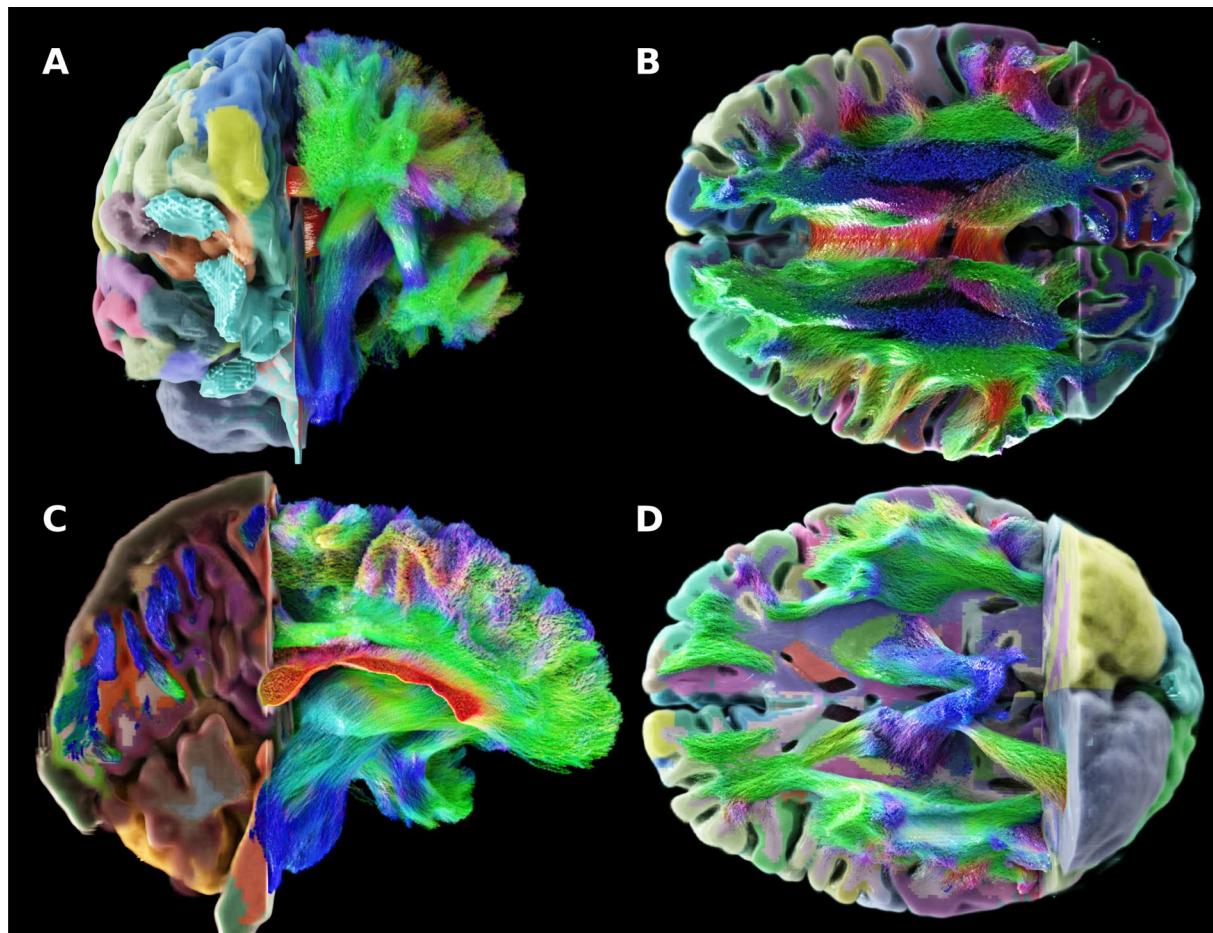


FIGURE 10 | Cinematic renderings of wholebrain tractography obtained using TractSeg, showing how they intersect with the cortical parcellation from FreeSurfer. The combination of tractography and functional parcellation allows estimation of whole-brain connectivity (Planchuelo-Gomez 2024). Coronal 3D views such as (A) highlight that corticospinal structure is inherently retrieved when using large FOV 3D imaging. (D) shows even cerebellar white matter pathways are mapped despite their relatively low diffusion anisotropy. (B) shows detailed depiction of association, projection and commissural pathways, while (C) clearly demonstrates the separation of the cingulum from the callosum.

otherwise technical and slow process; however, the encouraging results shown must be validated against high field data in a range of populations before this approach can be endorsed. Visual inspection suggests a high degree of correspondence with more traditional approaches to tractography like ACT, with the prior information that tracts are large coherent bundles clearly acting to regularise otherwise diffuse tracking. How these approaches may bias results, though, is as yet untested.

The role of AI in obtaining high quality tractography is significant, with these priors occurring in two stages: the first in bundle and endpoint segmentations, and the second in tract orientation mapping. Bundle segmentation restricts tracts to begin and end in particular regions, and the impact of this prior is visible comparing Figure 6 with Figure 4B,D,F. In Figure 4, manual ROI selection (itself a prior) correctly recovers the curved trajectory of the arcuate (B) and longitudinal fasciculus. These two features are similarly visible in Figure 6, though as these are derived from endpoint segmentation-filtered ACT rather than manual ROI ACT, they show more dispersion as they additionally select fibres in more cortical regions, away from the main trunk of the bundle.

Tract orientation mapping also substantially impacts tractography, with the strong priors asserted by TractSeg serving to regularise incoherency in the vector fields which embody tracts (Figure 9). This can be seen to only slightly deflect orientation vectors within the central trunk of larger tracts, with angular corrections appearing more aggressive in more peripheral and more cortical regions of tracts. As TractSeg is also able to use higher-order spherical harmonic peaks, it is difficult to extend this analysis to finer tracts, as it is less clear which peak (first, second or third largest) should be used to compare against.

Overall, it appears that for segmentation of major tracts, manual ROI approaches appear satisfactory; however, the use of AI-assisted approaches both simplifies and standardises segmentation while regularising inconsistencies in apparent tract orientation arising from noise or uncorrected artifacts. Despite promise, the impact of these DL approaches must be carefully and quantitatively evaluated for bias and accuracy before being incorporated in studies at scale.

While the cinematic rendering shown provides improved visual clarity and disambiguates complex topologies, it is important to note that we are still visualising the same interpolated

streamlines that have been presented in tractography studies since the late 1990s. The underlying data remain unchanged; what differs here is only the rendering style, which improves interpretability but does not alter or enhance the tractography estimates themselves.

The low inherent SNR and B_0 homogeneity of ULF-MRI requires the use of multi-shot imaging methods. DW-FSE provides interpretable, relatively robust diffusion encoded imaging with long echo trains providing many measurements per excitation. T_2 attenuation, which occurs during long readouts results in significant blurring arising from the reduced echo amplitude at the k-space periphery. This results in an echo-train length (ETL) dependant point-spread function (Zhou et al. 1993), which, alongside the low spatial resolution, prevents the recovery of smaller tracts or those in dense WM regions.

While diagnostic ULF-DWIs typically use high-resolution in-plane and thick slices, diffusion tractography instead benefits from isotropic resolution so as not to produce biased results in certain directions (Oouchi et al. 2007). Reduction of the in-plane resolution to make voxels isotropic benefits k-space SNR, as this is higher at more central k -space locations than more peripheral ones. This facilitates the use of longer echo trains than might otherwise be used for diagnostic sequences, and this work found extended ETLs significantly beyond those typical at ULF conferred substantial benefit to SNR-per-unit-time.

Corrections are demonstrated that address some of the issues that prevent immediate use of data for quantitative DWI and identify other issues that remain to be addressed. Significant work is still required to address some of the defects observed in the data, particularly spatial variation from gradient nonuniformity and static field gradients. Expectations of encoding uniformity that are typical at high field are shown to be erroneous in DW-ULF, and current tools for signal analysis fail to address this concern.

In conclusion, this study represents the first successful demonstration of anatomically faithful reconstructions of white matter pathways at ultra-low (64mT) field strength. With the rapid proliferation of low-field MRI systems in low- and middle-income countries (LMICs), this breakthrough has the potential to revolutionize neuroimaging, enabling the study of white matter architecture in regions where it was previously inaccessible. Having established the feasibility of tractography at ultra-low field strengths, future research will focus on enhancing the robustness of microstructural parameter quantification (tractometry) within these tracts. This is where we foresee the greatest impact, as it will enable truly democratized access to advanced neuroimaging techniques, facilitating broad applications in both research and clinical settings.

Acknowledgements

This work was funded in part by: The Gates Foundation under the UNITY programme (INV-047888); Wellcome LEAP (under the 1kD programme); A Wellcome Trust Strategic Award (104943/Z/14/Z); Wellcome Discovery Awards 227882/Z/23/Z and 317797/Z/24/Z. For

the purpose of open access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission. Steven Williams would like to thank the National Institute for Health and Care Research (NIHR) Maudsley Biomedical Research Centre (BRC) for ongoing support of our neuroimaging research. We thank Klaus Engel and Kaloian Petkov of Siemens Healthineers AG, Germany for producing the cinematic renderings in Figure 10.

Funding

This work was supported by the Bill and Melinda Gates Foundation, INV-047888; Wellcome Trust, 104943/Z/14/Z, 227882/Z/23/Z, 317797/Z/24/Z.

Conflicts of Interest

Francesco Padormo, Rui Teixeira and Rafael O'Halloran are employees of Hyperfine Inc. Klaus Engel and Kaloian Petkov are employees of Siemens Healthineers AG, Germany. All other authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

Andersson, J. L., M. S. Graham, E. Zsoldos, and S. N. Sotiroopoulos. 2016. "Incorporating Outlier Detection and Replacement Into a Non-Parametric Framework for Movement and Distortion Correction of Diffusion MR Images." *NeuroImage* 141: 556–572.

Andersson, J. L., S. Skare, and J. Ashburner. 2003. "How to Correct Susceptibility Distortions in Spinecho Echo-Planar Images: Application to Diffusion Tensor Imaging." *NeuroImage* 20: 870–888. <https://pubmed.ncbi.nlm.nih.gov/14568458/>.

Andersson, J. L., and S. N. Sotiroopoulos. 2016. "An Integrated Approach to Correction for Off-Resonance Effects and Subject Movement in Diffusion MR Imaging." *NeuroImage* 125: 1063–1078.

Balaji, S., N. Wiley, A. Dvorak, et al. 2025. "Magnetization Transfer Imaging Using Non-Balanced SSFP at Ultra-Low Field." *Magnetic Resonance in Medicine* 94: 602–614.

Cahn, B. A., J. T. Shah, H. A. Dyvorne, et al. 2020. "Abstract 57: Deployment of Portable, Bedside, Low-Field Magnetic Resonance Imaging for Evaluation of Stroke Patients." *Stroke* 51, no. Suppl-1. <https://doi.org/10.1161/STR.51.suppl.1.57>.

Calamuneri, A., A. Arrigo, E. Mormina, et al. 2018. "White Matter Tissue Quantification at Low b-Values Within Constrained Spherical Deconvolution Framework." *Frontiers in Neurology* 9: 386276. www.frontiersin.org.

Canales-Rodríguez, E. J., A. Daducci, S. N. Sotiroopoulos, et al. 2015. "Spherical Deconvolution of Multichannel Diffusion MRI Data With Non-Gaussian Noise Models and Spatial Regularization." *PLoS One* 10: e0138910. <https://doi.org/10.1371/journal.pone.0138910>.

Catani, M., R. J. Howard, S. Pajevic, and D. K. Jones. 2002. "Virtual In Vivo Interactive Dissection of White Matter Fasciculi in the Human Brain." *NeuroImage* 17: 77–94. <https://pubmed.ncbi.nlm.nih.gov/12482069/>.

Chang, L. C., D. K. Jones, and C. Pierpaoli. 2005. "RESTORE: Robust Estimation of Tensors by Outlier Rejection." *Magnetic Resonance in Medicine* 53: 1088–1095. <https://doi.org/10.1016/j.neuroimage.2019.116137>.

Chang, L. C., L. Walker, and C. Pierpaoli. 2012. "Informed RESTORE: A Method for Robust Estimation of Diffusion Tensor From Low

Redundancy Datasets in the Presence of Physiological Noise Artifacts." *Magnetic Resonance in Medicine* 68: 1654–1663. <https://doi.org/10.1002/mrm.24173>.

Chung, S. W., Y. Lu, and R. G. Henry. 2006. "Comparison of Bootstrap Approaches for Estimation of Uncertainties of DTI Parameters." *NeuroImage* 33: 531–541.

Collins, D. L., P. Neelin, T. M. Peters, and A. C. Evans. 1994. "Automatic 3D Intersubject Registration of MR Volumetric Data in Standardized Talairach Space." *Journal of Computer Assisted Tomography* 18: 192–205. <https://www.scienceopen.com/document?vid=25d2ff94-4ac7-4b7c-8afcba7b2e6cf273>.

Comaniciu, D., K. Engel, B. Georgescu, and T. Mansi. 2016. "Shaping the Future Through Innovations: From Medical Imaging to Precision Medicine." *Medical Image Analysis* 33: 19–26.

Conturo, T. E., N. F. Lori, T. S. Cull, et al. 1999. "Tracking Neuronal Fiber Pathways in the Living Human Brain." *Proceedings of the National Academy of Sciences of the United States of America* 96: 10422–10427. <https://pubmed.ncbi.nlm.nih.gov/10468624/>.

Dale, A., B. Fischl, and M. I. Sereno. 1999. "Cortical Surface-Based Analysis: I. Segmentation and Surface Reconstruction." *NeuroImage* 9: 179–194.

Dappa, E., K. Higashigaito, J. Fornaro, et al. 2016. "Cinematic Rendering – An Alternative to Volume Rendering for 3D Computed Tomography Imaging." *Insights Into Imaging* 7: 849–856. <https://doi.org/10.1007/s13244-016-0518-1>.

Dell'Acqua, F., M. Descoteaux, and A. Leemans. 2024. "Handbook of Diffusion MR Tractography: Imaging Methods, Biophysical Models, Algorithms and Applications." In *Handbook of Diffusion MR Tractography: Imaging Methods, Biophysical Models, Algorithms and Applications*, 1–727. Academic Press.

Dhollander, T., D. Raffelt, and A. Connelly. 2016. "Unsupervised 3-Tissue Response Function Estimation From Single-Shell or Multi-Shell Diffusion MR Data Without a Co-Registered T1 Image in ISMRM Workshop on Breaking the Barriers of Diffusion MRI." https://www.researchgate.net/publication/307863133_Uncsupervised_3-tissue_response_function_estimation_from_single-shell_or_multi-shell_diffusion_MR_data_without_a_co-registered_T1_image.

Ding, Y. 2024. "Diffusion Tensor Imaging at 0.05 T in Proceedings of the International Society for Magnetic Resonance in Medicine 2024."

Fedorov, A., R. Beichel, J. Kalpathy-Cramer, et al. 2012. "3D Slicer as an Image Computing Platform for the Quantitative Imaging Network." *Magnetic Resonance Imaging* 30: 1323–1341.

Fischl, B., A. Liu, and A. M. Dale. 2001. "Automated Manifold Surgery: Constructing Geometrically Accurate and Topologically Correct Models of the Human Cerebral Cortex." *IEEE Transactions on Medical Imaging* 20: 70–80.

Fischl, B., D. H. Salat, E. Busa, et al. 2002. "Whole Brain Segmentation: Automated Labeling of Neuroanatomical Structures in the Human Brain." *Neuron* 33: 341–355.

Fischl, B., D. H. Salat, A. J. W. van der Kouwe, et al. 2004. "Sequence-Independent Segmentation of Magnetic Resonance Images." *NeuroImage Mathematics in Brain Imaging* 23: S69–S84.

Fischl, B., M. I. Sereno, and A. Dale. 1999. "Cortical Surface-Based Analysis: II: Inflation, Flattening, and a Surface-Based Coordinate System." *NeuroImage* 9: 195–207.

Garyfallidis, E., M. Brett, B. Amirbekian, et al. 2014. "Dipy, a Library for the Analysis of Diffusion MRI Data." *Frontiers in Neuroinformatics* 8: 71423.

Gholam, J., Á. Planchuelo-Gómez, J. Ametepe, et al. 2024. "Resolving Complex White Matter Architecture in the Human Brain at 64 mT." In *Proceedings of the International Society for Magnetic Resonance in Medicine*. International Society for Magnetic Resonance in Medicine.

Griswold, M. A., D. Walsh, R. M. Heidemann, A. Haase, and P. M. Jakob. 2002. "The Use of an Adaptive Reconstruction for Array Coil Sensitivity Mapping and Intensity Normalization in Proceedings of the International Society for Magnetic Resonance in Medicine 2002."

Iglesias, J. E., B. Billot, Y. Balbastre, et al. 2021. "Joint Super-Resolution and Synthesis of 1 Mm Isotropic MP-RAGE Volumes From Clinical MRI Exams With Scans of Different Orientation, Resolution and Contrast." *NeuroImage* 237: 118206. <https://www.sciencedirect.com/science/article/pii/S1053811921004833>.

Iglesias, J. E., B. Billot, Y. Balbastre, et al. 2023. "SynthSR: A Public AI Tool to Turn Heterogeneous Clinical Brain Scans Into High-Resolution T1-Weighted Images for 3D Morphometry." *Science Advances* 9, no. 5: eadd3607. <https://doi.org/10.1126/sciadv.add3607>.

Iglesias, J. E., R. Schleicher, S. Laguna, et al. 2022. "Quantitative Brain Morphometry of Portable Low-Field-Strength MRI Using Super-Resolution Machine Learning." *Radiology* 306, no. 3: e220522. <https://doi.org/10.1148/radiol.220522>.

Jeurissen, B., J. D. Tournier, T. Dhollander, A. Connelly, and J. Sijbers. 2014. "Multi-Tissue Constrained Spherical Deconvolution for Improved Analysis of Multi-Shell Diffusion MRI Data." *NeuroImage* 103: 411–426. <https://pubmed.ncbi.nlm.nih.gov/25109526/>.

Jones, D. K. 2010. "Optimal Approaches to Diffusion MRI Acquisition." In *Diffusion MRI*, 250–271. Oxford University Press. <https://academic.oup.com/book/24921/chapter/188777735>.

Jones, D. K., and P. J. Basser. 2004. "Squashing Peanuts and Smashing Pumpkins": How Noise Distorts Diffusion-Weighted MR Data." *Magnetic Resonance in Medicine* 52: 979–993. <https://doi.org/10.1002/mrm.20283>.

Kellner, E., B. Dhital, V. G. Kiselev, and M. Reisert. 2016. "Gibbs-Ringing Artifact Removal Based on Local Subvoxel-Shifts." *Magnetic Resonance in Medicine* 76: 1574–1581. <https://pubmed.ncbi.nlm.nih.gov/26745823/>.

Koay, C. G., and P. J. Basser. 2006. "Analytically Exact Correction Scheme for Signal Extraction From Noisy Magnitude MR Signals." *Journal of Magnetic Resonance* 179: 317–322.

Leemans, A., and D. K. Jones. 2009. "The B-Matrix Must Be Rotated When Correcting for Subject Motion in DTI Data." *Magnetic Resonance in Medicine* 61: 1336–1349. <https://pubmed.ncbi.nlm.nih.gov/19319973/>.

Lena, B. 2025. "Reproducibility and Repeatability of Rapid T1 Mapping of Brain Tissues at 64 mT: A Multicenter Study in Proceedings of the International Society for Magnetic Resonance in Medicine 2025."

O'Halloran, R. 2022. "Diffusion-Weighted Imaging at 0.064 T in Proceedings of the International Society for Magnetic Resonance in Medicine 2022."

Ouchi, H., K. Yamada, K. Sakai, et al. 2007. "Diffusion Anisotropy Measurement of Brain White Matter Is Affected by Voxel Size: Underestimation Occurs in Areas With Crossing Fibers." *American Journal of Neuroradiology* 28: 1102. <https://pmc.ncbi.nlm.nih.gov/articles/PMC8134156/>.

Park, S., J. M. Witten, and K. J. Myers. 2009. "Singular Vectors of a Linear Imaging System as Efficient Channels for the Bayesian Ideal Observer." *IEEE Transactions on Medical Imaging* 28: 657–668.

Planchuelo-Gomez, A. 2024. "Connectomics at 64 mT in Proceedings of the International Society for Magnetic Resonance in Medicine 2024."

Plumley, A., F. Padormo, M. Cercignani, et al. 2022. "Tensors and Tracts at 64 mT in ISMRM Workshop on Diffusion MRI: From Research to Clinic."

Schroeder, W., K. Martin, and B. Lorensen. 2006. "The Visualization Toolkit, 4th ed."

Sinha, A., and K. N. Chaudhury. 2024. “FISTA Iterates Converge Linearly for Denoiser-Driven Regularization.” <http://arxiv.org/abs/2411.10808>.

Smith, R. E., J. D. Tournier, F. Calamante, and A. Connelly. 2012. “Anatomically-Constrained Tractography: Improved Diffusion MRI Streamlines Tractography Through Effective Use of Anatomical Information.” *NeuroImage* 62: 1924–1938.

Smith, S. M., M. Jenkinson, M. W. Woolrich, et al. 2004. “Advances in Functional and Structural MR Image Analysis and Implementation as FSL.” *NeuroImage* 23, no. Suppl 1: S208–S219. <https://pubmed.ncbi.nlm.nih.gov/15501092/>.

Sorby-Adams, A., J. Kirsch, J. Guo, et al. 2024. “Abstract 128: Multi-Direction Diffusion Weighted Imaging on Portable, Low-Field Magnetic Resonance Imaging.” *Stroke* 55, no. Suppl-1. <https://doi.org/10.1161/str.55.suppl.1.128>.

Tax, C. M., M. Bastiani, J. Veraart, E. Garyfallidis, and M. O. Irfanoglu. 2022. “What’s New and What’s Next in Diffusion MRI Preprocessing.” *NeuroImage* 249: 118830. <https://pubmed.ncbi.nlm.nih.gov/34965454/>.

Tournier, J. D., F. Calamante, and A. Connelly. 2013. “Determination of the Appropriate b Value and Number of Gradient Directions for High-Angular-Resolution Diffusion-Weighted Imaging.” *NMR in Biomedicine* 26: 1775–1786.

Tournier, J. D., R. Smith, D. Raffelt, et al. 2019. “MRtrix3: A Fast, Flexible and Open Software Framework for Medical Image Processing and Visualisation.” *NeuroImage* 202: 116137.

Tournier, J. D., C.-H. Yeh, F. Calamante, et al. 2008. “Resolving Crossing Fibres Using Constrained Spherical Deconvolution: Validation Using Diffusion-Weighted Imaging Phantom Data.” *NeuroImage* 42: 617–625.

Tustison, N. J., B. B. Avants, P. A. Cook, et al. 2010. “N4ITK: Improved N3 Bias Correction.” *IEEE Transactions on Medical Imaging* 29: 1310. <https://pmc.ncbi.nlm.nih.gov/articles/PMC3071855/>.

Veraart, J., D. S. Novikov, D. Christiaens, B. Ades-aron, J. Sijbers, and E. Fieremans. 2016. “Denoising of Diffusion MRI Using Random Matrix Theory.” *NeuroImage* 142: 394–406. <https://pmc.ncbi.nlm.nih.gov/articles/PMC5159209/>.

Vos, S. B., C. M. W. Tax, P. R. Luijten, et al. 2017. “The Importance of Correcting for Signal Drift in Diffusion MRI.” *Magnetic Resonance in Medicine* 77: 285–299.

Walsh, D. O., A. F. Gmitro, and M. W. Marcellin. 2000. “Adaptive Reconstruction of Phased Array MR Imagery in Proceedings of the International Society for Magnetic Resonance in Medicine 2000.”

Wasserthal, J., P. Neher, and K. H. Maier-Hein. 2018a. “TractSeg – Fast and Accurate White Matter Tract Segmentation.” *NeuroImage* 183: 239–253.

Wasserthal, J., P. F. Neher, and K. H. Maier-Hein. 2018b. “Tract Orientation Mapping for Bundle-Specific Tractography.” *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 11072: 36–44.

Yendiki, A., K. Koldewyn, S. Kakunoori, N. Kanwisher, and B. Fischl. 2014. “Spurious Group Differences due to Head Motion in a Diffusion MRI Study.” *NeuroImage* 88: 79–90.

Zhou, X., Z. P. Liang, G. P. Cofer, et al. 1993. “Reduction of Ringing and Blurring Artifacts in Fast Spin-Echo Imaging.” *Journal of Magnetic Resonance Imaging* 3: 803–807.

Appendix A

Presuming the spatially dependent inhomogeneous field $B_0(\mathbf{r})$ contributes a significant spatially dependent gradient:

$$\mathbf{G}(\mathbf{r}) = \frac{\partial B_0(\mathbf{r})}{\partial x} + \frac{\partial B_0(\mathbf{r})}{\partial y} + \frac{\partial B_0(\mathbf{r})}{\partial z} \quad (\text{A1})$$

we may consider an expression for a perturbation to the effective b-value and direction of the gradient during diffusion encoding. The effective encoding direction is given by

$$\mathbf{G}_e(\mathbf{r}, t) = \mathbf{G}_D(t) + \mathbf{G}(\mathbf{r}) \quad (\text{A2})$$

where $\mathbf{G}_D(t)$ is the prescribed, time-varying diffusion gradient and $\mathbf{G}_e(\mathbf{r}, t)$ is the effective diffusion gradient. Assuming encoding is dominated by that which occurs while $\mathbf{G}_D(t)$ plays, we can readily model the contribution to individual terms of the spatially dependent b-matrix $b_{ij}(\mathbf{r})$. We describe the spatially varying phase encoding given a gradient vector $\mathbf{G}(\mathbf{r}, t)$ as $\mathbf{F}(\mathbf{r}, t)$:

$$\mathbf{F}(\mathbf{r}, t) = \int_0^{TE} \mathbf{G}_e(\mathbf{r}, t') dt' \quad (\text{A3})$$

which is used to compute the pairwise components of the b-matrix in the scanner’s frame of reference:

$$b_{ij}(\mathbf{r}) = \int_0^{TE} \mathbf{F}_i(\mathbf{r}, t) \mathbf{F}_j(\mathbf{r}, t) dt \quad (\text{A4})$$

We can see from this that the effect of a static field gradient on typical Stejskal-Tanner encoding is to both rotate and scale the encoding gradient. Rotation effects are relatively small for large ratios of G_D to $\mathbf{G}(\mathbf{r})$, even for orthogonal gradients. A large inhomogeneity field with the static gradient orthogonal to the prescribed gradient, with a length ratio of 20:1 would produce a rotation of the effective encoding of 2.9°. Comparatively, scaling effects on the interpretation of diffusion encoding are comparatively significant and more pronounced when the prescribed and static gradients are colinear. Using the \mathbf{G}^2 scaling implied by Equation (A4), the true encoding magnitude will vary with the quadratic of Equation (A2). The effective scaling of the encoding magnitude from that prescribed may be considered as a scalar a multiplying the b-value that is

$$b_{\text{effective}} = ab \quad (\text{A5})$$

And this expression is used to modify the Stejskal-Tanner signal equation:

$$S' = S_0 \exp(-abD) \quad \text{hence} \quad S = \exp\left(\log\left(\frac{S'}{S_0}\right)\left(1 - \frac{1}{a}\right)\right) \quad (\text{A6})$$

it is possible to modify the measured signal given knowledge of S_0 and the scaling factor a , which scales approximately with the ratio $\left(1 + \frac{G_D}{\mathbf{G}(\mathbf{r})}\right)^2$. The actual impact of this effect then is highly dependent on both the diffusion encoding gradient used and the severity of the inhomogeneity.

Though the correction demonstrated is suitable for magnitude correction of encoding errors, we may alternatively solve for a spatially dependent b-matrix, as is used for magnetic field gradient inhomogeneity correction. Though this method is more general, it does not readily extend to CSD where a single response function representative of all voxels is required.