NODDI: practical *in vivo* neurite orientation dispersion and density imaging of the human brain

Hui Zhang[a,*], Torben Schneider[b], Claudia A Wheeler-Kingshott[b], Daniel C Alexander[a]

[a] *Department of Computer Science & Centre for Medical Image Computing, University College London, UK*
[b] *NMR Research Unit, Department of Neuroinflammation, UCL Institute of Neurology, University College London, UK*

Abstract

This paper introduces neurite orientation dispersion and density imaging (NODDI), a practical diffusion MRI technique for estimating the microstructural complexity of dendrites and axons *in vivo* on clinical MRI scanners. Such indices of neurites relate more directly to and provide more specific markers of brain tissue microstructure than standard indices from diffusion tensor imaging, such as fractional anisotropy (FA). Mapping these indices over the whole brain on clinical scanners presents new opportunities for understanding brain development and disorders. The proposed technique enables such mapping by combining a three-compartment tissue model with a two-shell high-angular-resolution diffusion imaging (HARDI) protocol optimized for clinical feasibility. An index of orientation dispersion is defined to characterize angular variation of neurites. We evaluate the method both in simulation and on a live human brain using a clinical 3T scanner. Results demonstrate that NODDI provides sensible neurite density and orientation dispersion estimates, thereby disentangling two key contributing factors to FA and enabling the analysis of each factor individually. We additionally show that while orientation dispersion can be estimated with just a single HARDI shell, neurite density requires at least two shells and can be estimated more accurately with the optimized two-shell protocol than with alternative two-shell protocols. The optimized protocol takes about 30 minutes to acquire, making it feasible for inclusion in a typical clinical setting. We further show that sampling fewer orientations in each shell can reduce the acquisition time to just 10 minutes with minimal impact on the accuracy of the estimates. This demonstrates the feasibility of NODDI even for the most time-sensitive clinical applications, such as neonatal and dementia imaging.

Keywords: Neurite density, Orientation dispersion, Diffusion MRI

1. Introduction

Dendrites and axons, known collectively as neurites, are projections of neurons. They are the cellular building blocks of the computational circuitry of the brain. Quantifying neurite morphology in terms of its density and orientation distribution provides a window into the structural basis of brain function both in normal populations and in populations with brain disorders. For example, the branching complexity of the dendritic trees, measured in terms of dendritic density, reflects the nature of their computation and hence their function: the areas of the cortex with less complex dendritic structures engage in the early stages of information processing while the regions with more complex dendritic structures participate in the later stages of processing (Jacobs et al., 2001). Neurite morphology is also a key marker of brain development and aging. An increase in the dispersion of neurite orientation distribution is associated with brain development (Conel, 1939), whereas a reduction in the dendritic density is linked with the aging of the brain (Jacobs et al., 1997). Changes in neurite morphology are additionally implicated in

*Corresponding author
Email address: g.zhang@cs.ucl.ac.uk (Hui Zhang)

numerous neurological disorders, including multiple sclerosis (Evanglou et al., 2000), amyotrophic lateral sclerosis (Bruijn et al., 2004), and Alzheimer's disease (Paula-Barbosa et al., 1980). However, due to its reliance on scarcely available postmortem tissue samples, the quantitative analysis of neurite morphology, despite its importance, is not widely applied. The development of a non-invasive imaging-based solution holds the key to realize such quantification *in vivo*.

Diffusion magnetic resonance imaging (MRI) provides unique insight into tissue microstructure and is arguably the most promising candidate for *in vivo* quantification of neurite morphology. It works by sensitizing MRI measurements to the displacement pattern of water molecules undergoing diffusion. As the water displacement pattern is influenced by tissue microstructure, by measuring this displacement pattern, diffusion MRI is able to distinguish different microstructural environments. In the case of neuronal tissues, during the typical time scale of a diffusion MRI experiment, two kinds of microstructural environments can be identified, which are characterized by either hindered or restricted diffusion (Assaf and Cohen, 2000). Hindered diffusion refers to the diffusion of water with a Gaussian displacement pattern. It characterizes the water in the extra-cellular space defined by cellular membranes of somas and glial cells. Restricted diffusion refers to the diffusion of water in restricted geometries. It is characterized by a non-Gaussian pattern of displacement and describes the water in the intra-cellular space bounded for example by axonal or dendritic membranes. The differentation of intra- and extra-cellular water forms the basis of measuring neurite morphology via diffusion MRI.

Currently, the standard clinical diffusion MRI technique is diffusion tensor imaging (DTI) (Basser et al., 1994). This technique provides sensitivity to tissue microstructure but lacks specificity for individual tissue microstructure features (Pierpaoli et al., 1996). DTI provides simple markers, such as mean diffusivity (MD) and fractional anisotropy (FA), that are widely used as surrogate measures of microstructural tissue change during normal brain development and aging, or during the onset and progression of neurological disorders (See, e.g., Salat et al. (2009) and Bodini and Ciccarelli (2009) for reviews). However, despite their sensitivity, these markers are inherently non-specific (Pierpaoli et al., 1996). For instance, the observation of a reduction in FA may be caused by a reduction in neurite density, an increase in the dispersion of neurite orientation distribution, as well as various other tissue microstructural changes (Beaulieu, 2009). Hence, a change in these statistics may not be attributed to specific changes in tissue microstructure.

Towards *in vivo* quantification of neurite morphology, the recent trend in diffusion MRI is in developing more advanced techniques that can measure tissue microstructure features directly (See Assaf and Cohen (2009) for a review). A particularly successful approach, pioneered by Stanisz et al. (1997), is the model-based strategy in which a geometric model of the microstructure of interest predicts the MR signal from water diffusion within. The authors propose a model of white matter microstructure that consists of individual compartments for glial cells, axons, and extra-cellular space. The glial and axon compartments have restricted diffusion; extra-cellular diffusion is hindered with apparent diffusivities calculated via a tortuosity model. The model allows exchange between the intra-cellular and extra-cellular compartments.

Subsequent white matter models include the ball-and-stick model (Behrens et al., 2003), which represents the intra-cellular compartment as cylinders of zero radius and extra-cellular diffusion as isotropic and unrestricted. The composite hindered and restricted water diffusion (CHARMED) model (Assaf et al., 2004; Assaf and Basser, 2005) represents the intra-cellular compartment as impermeable parallel cylinders with a gamma distribution of radii. The signal for the extra-cellular compartment comes from an anisotropic diffusion tensor model. In the original model, the distribution of radii is fixed to a biologically plausible distribution, but subsequent work (Assaf et al., 2008) fits these parameters. Barazany et al. (2009) add a free-water compartment necessary for *in vivo* imaging data. Alexander (2008) reduces the CHARMED model to a single radius, and subsequently Alexander et al. (2010) include tortuosity models and isotropically restricted compartments, in a similar way to Stanisz et al. (1997), and a free-water compartment as in Barazany et al. (2009), to obtain the minimal model of white matter diffusion (MMWMD). Most recently, (Panagiotaki et al., 2012) construct a taxonomy of compartment models for white matter including those above and a range of intermediate and additional compartment combinations. They compare them with each other and with various multi-exponential models using fixed-brain data to demonstrate the need for both a restricted axonal compartment and an isotropically restricted compartment as in Stanisz et al. (1997); Alexander et al. (2010).

Assaf and Basser (2005) demonstrate for the first time that the CHARMED model can provide sensible maps of the volume fraction of intra-cellular space, the axon density, in *in vivo* human brain imaging on a clinical MRI scanner. However, by representing axons as parallel cylinders, models such as ball-and-stick (Behrens et al., 2003), CHARMED, MMWMD, and the entire hierarchy of compartment models in Panagiotaki et al. (2012) can not recover the effect of axonal-orientation dispersion due to bending and fanning of axon bundles widespread throughout the brain (House and Pansky, 1960; Bürgel et al., 2006). By relaxing this constraint, more recent models (Kaden et al., 2007; Zhang et al., 2011; Sotiropoulos et al., 2012) support a more realistic description of white matter beyond the most coherently-oriented structures, such as the corpus callosum, and provide an estimate of orientation dispersion.

Going beyond the modeling of white matter, Jespersen et al. (2007) propose an analytic model of neurites that support the modeling of both gray and white matter. Using a truncated spherical harmonic series, the neurite model approximates an arbitrary orientation distribution of dendrites and axons, which is essential for modeling both low-to-moderately dispersed axons in white matter and highly dispersed dendritic trees in gray matter. Using data from *ex vivo* imaging of a baboon brain sample, the authors demonstrate for the first time that both neurite density and its orientation distribution can be quantified using diffusion MRI. However, the imaging protocol consists of 153 diffusion-weighted images spread over 17 b-values with the largest equal to 15,000 s/mm^2, making it impractical for clinical translation.

Despite the lack of clinically feasible imaging protocols, emerging evidence suggests that, in both gray and white matter, neurite morphology determined from diffusion MRI is comparable to independent measures derived from histology. Jespersen et al. (2010) show that neurite density estimates, determined using the model in Jespersen et al. (2007), correlate more strongly with both optical myelin staining intensity and stereological estimation of neurite density using electron microscopy than DTI-derived markers. More recently, Jespersen et al. (2012) demonstrate that neurite orientation distributions derived from diffusion MRI show excellent agreement to those quantified using a quantitative Golgi analysis. These recent findings are extremely encouraging and motivate the current work.

The aim of this work is to develop a clinically feasible technique for *in vivo* neurite orientation dispersion and density imaging, which we refer to hereafter as NODDI. Our approach is to first choose a model that is sufficiently simple, yet complex enough to capture the key features of neurite morphology, then identify the optimal acquisition protocol for such a model under scanner hardware and acquisition time constraints typical in a clinical setting. Specifically, NODDI adapts the orientation-dispersed cylinder model in Zhang et al. (2011) to estimate only neurite density and orientation dispersion. The acquisition protocol is determined using the experiment design optimization in Alexander (2008) under an acquisition time constraint of 30 minutes. Using both synthetic and *in vivo* human brain data, we assess the performance of the optimized protocol, in terms of the accuracy and precision of its microstructure parameter estimates, against alternative protocols.

The rest of the paper is organized as follows: Section 2 describes the NODDI tissue model, protocol optimization, data acquisition, model fitting, and preprocessing; Section 3 gives the experimental design and results; Section 4 summarizes the contribution and discusses future work.

2. Materials and methods

This section specifies the NODDI tissue model and defines the orientation dispersion index, which is NODDI's summary statistic for quantifying angular variation of neurite orientation. It then details the protocol optimization for this model, the resulting NODDI protocol, and the alternative protocols for comparison. Finally, it describes the acquisition of the *in vivo* imaging data, the synthesis of the simulated data, the model-fitting procedures, and the preprocessing of the imaging data.

2.1. NODDI tissue model

NODDI adopts a tissue model that distinguishes three types of microstructural environment: intra-cellular, extra-cellular, and CSF compartments. Each affects water diffusion within the environment in an

unique way (Le Bihan, 1995) and gives rise to a separate normalized MR signal. The full normalized signal A can be written as

$$A = (1 - \nu_{iso})(\nu_{ic} A_{ic} + (1 - \nu_{ic}) A_{ec}) + \nu_{iso} A_{iso}\,, \tag{1}$$

where A_{ic} and ν_{ic} are the normalized signal and volume fraction of the intra-cellular compartment; A_{ec} is the normalized signal of the extra-cellular compartment; A_{iso} and ν_{iso} are the normalized signal and volume fraction of the CSF compartment. The tissue and signal model for each compartment is detailed below.

2.1.1. Intra-cellular model

The intra-cellular compartment refers to the space bounded by the membrane of neurites. We model this space as a set of *sticks*, i.e., cylinders of zero radius, to capture the highly restricted nature of diffusion perpendicular to neurites and unhindered diffusion along them (Behrens et al., 2003; Panagiotaki et al., 2012; Sotiropoulos et al., 2012). The orientation distribution of sticks can range from highly parallel to highly dispersed. This models the full spectrum of neurite orientation patterns observed in brain tissue that include: 1) highly coherently oriented white matter structures, such as the corpus callosum; 2) white matter structures composed of bending and fanning axons, such as the centrum semiovale; 3) the cerebral cortex and subcortical gray matter structures characterized by sprawling dendritic processes in all directions.

The normalized signal, A_{ic}, adopts the orientation-dispersed cylinder model in Zhang et al. (2011) but is simplified for sticks, such that

$$A_{ic} = \int_{\mathbb{S}^2} f(\mathbf{n})\, e^{-b d_{\parallel} (\mathbf{q}\cdot\mathbf{n})^2}\, d\mathbf{n}\,, \tag{2}$$

where $\mathbf{q}$ and b are the gradient direction and b-value of diffusion-weighting, respectively; $f(\mathbf{n})d\mathbf{n}$ gives the probability of finding sticks along orientation $\mathbf{n}$; $e^{-b d_{\parallel} (\mathbf{q}\cdot\mathbf{n})^2}$ gives the signal attenuation due to unhindered diffusion along a stick with intrinsic diffusivity $d_{\parallel}$ and orientation $\mathbf{n}$.

As in Zhang et al. (2011), the orientation distribution function $f : \mathbb{S}^2 \mapsto \mathbb{R}$ is modeled with a Watson distribution:

$$f(\mathbf{n}) = M\left(\frac{1}{2}, \frac{3}{2}, \kappa\right)^{-1} e^{\kappa(\boldsymbol{\mu}\cdot\mathbf{n})^2}\,, \tag{3}$$

where M is a confluent hypergeometric function, $\boldsymbol{\mu}$ is the mean orientation, and κ is the concentration parameter that measures the extent of orientation dispersion about $\boldsymbol{\mu}$.

The Watson distribution is chosen because it is the simplest orientation distribution that can capture the dispersion in orientations (Mardia and Jupp, 1990). Compared to the truncated spherical harmonic series in Jespersen et al. (2007), our choice trades the generality in representing complex orientation distributions for the simplicity in capturing the essence of orientation dispersion. Furthermore, the Watson distribution provides a good representation not only for high orientation dispersion seen in gray matter but also for low orientation dispersion seen in the most coherent white matter. In contrast, the truncated spherical harmonics series, with the maximum order of spherical harmonics set typically to 2 or 4 (Jespersen et al., 2007), can not approximate the most coherent orientation distributions accurately (Zhang et al., 2011). Including higher order terms increases the number of model parameters significantly, which outweighs the gain in improved approximation.

2.1.2. Extra-cellular model

The extra-cellular compartment refers to the space around the neurites, which is occupied by various types of glial cells and, additionally in gray matter, cell bodies (somas). In this space, the diffusion of water molecules is hindered by the presence of neurites but not restricted, hence is modeled with simple (Gaussian) anisotropic diffusion.

The normalized signal, A_{ec}, again adopts the extra-cellular signal model of orientation-dispersed cylinders in Zhang et al. (2011), such that

$$\log A_{ec} = -b\, \mathbf{q}^T \left(\int_{\mathbb{S}^2} f(\mathbf{n}) D(\mathbf{n})\, d\mathbf{n} \right) \mathbf{q}\,, \tag{4}$$

where $D(\mathbf{n})$ is a cylindrically symmetric tensor with the principal direction of diffusion $\mathbf{n}$, diffusion coefficients $d_\parallel$ parallel to $\mathbf{n}$ and $d_\perp$ perpendicular to $\mathbf{n}$. The parallel diffusivity is the same as the intrinsic free diffusivity of the intra-cellular compartment; the perpendicular diffusivity is set with a simple tortuosity model (Szafer et al., 1995) as $d_\perp = d_\parallel(1 - \nu_{ic})$, where ν_{ic} is the intra-cellular volume fraction.

In the current implementation with the Watson distribution, the parallel and perpendicular diffusivities of the apparent extra-cellular diffusion tensor, denoted as $d'_\parallel$ and $d'_\perp$, take the following form (Zhang et al., 2011):

$$d'_\parallel = d_\parallel - d_\parallel \, \nu_{ic} \, (1 - \tau_1) \tag{5}$$

$$d'_\perp = d_\parallel - d_\parallel \, \nu_{ic} \, \frac{1 + \tau_1}{2} \, , \tag{6}$$

where, borrowing the definition from Jespersen et al. (2012),

$$\tau_1 = \int_0^1 u^2 e^{\kappa u^2} \, du \, / \int_0^1 e^{\kappa u^2} \, du \, , \tag{7}$$

which captures the effect of orientation dispersion on the apparent diffusivities. Written in terms of the Dawson's integral (Abramowitz and Stegun, 1972), $F(x) = \frac{1}{2}\sqrt{\pi} e^{-x^2} \texttt{erfi}(x)$, τ_1 is equal to

$$-\frac{1}{2\kappa} + \frac{1}{2F(\sqrt{\kappa})\sqrt{\kappa}} \, . \tag{8}$$

It varies from 1/3 for isotropically-dispersed orientations ($\kappa = 0$) to 1 for strictly parallel orientations ($\kappa = \infty$).

Eqns (5) and (6) demonstrate a significant difference between our extra-cellular model and the earlier models such as ball-and-stick (Behrens et al., 2003), CHARMED, and the neurite model in Jespersen et al. (2007). Unlike these earlier models which treat the apparent parallel and perpendicular diffusivities in the extra-cellular space as independent free parameters, our model expresses them in terms of the neurite morphology and the intrinsic diffusivity in a physically plausible way. In particular, these parameters are now determined not only by the neurite density, via the tortuosity model, but also by the orientation dispersion of neurites, via the powder average in Eqn (4). For instance, an increase in orientation dispersion will not only yield a reduction of the intra-cellular anisotropy via Eqn (2) but also lead to a reduction of the extra-cellular anisotropy via a reduction in the apparent extra-cellular parallel diffusivity and an increase in the corresponding apparent perpendicular diffusivity.

2.1.3. CSF compartment

The CSF compartment models the space occupied by cerebrospinal fluid and is modeled as isotropic Gaussian diffusion with diffusivity d_{iso}.

2.2. Orientation dispersion index

Here we redefine the *orientation dispersion index* first proposed in Zhang et al. (2011) as

$$\mathrm{OD} = \frac{2}{\pi} \arctan\left(1/\kappa\right) . \tag{9}$$

In Zhang et al. (2011), we defined the index simply as κ, which is less intuitive because it maps higher orientation dispersion to lower values. The new definition addresses this issue. Furthermore, it ranges from 0 to 1, making it more straightforward to visualize than κ which has an upper bound that is infinity. Fig. 1 illustrates the Watson distributions for a range of OD values.

2.3. Protocol optimization

For the NODDI tissue model, an optimized protocol is derived using the experiment design optimization procedure in Alexander (2008). The procedure determines the optimal protocol for a set of *a priori* model parameter settings and the hardware specification under an acquisition time constraint. We choose the set of *a priori* model parameters representative for both gray and white matter: intra-cellular volume fraction ν_{ic} = 0.3, 0.5, and 0.7, intrinsic free diffusivity $d_{\parallel} = 1.7\times10^{-3}$ mm^{2}s^{-1}, perpendicular extra-cellular diffusivity $d_{\perp} = 1.2\times10^{-3}$, 0.9×10^{-3}, and 0.5×10^{-3} mm^{2}s^{-1}, κ = 0.5, 2, 8, and 32. The hardware specification of the scanner used for the *in vivo* imaging demonstration is given in Section 2.4. The acquisition time limit is set to 30 minutes during which the scanner can acquire around 90 diffusion-weighted images with whole-brain coverage.

2.3.1. NODDI protocol

The optimization procedure suggests dividing the measurements into 2 HARDI shells with the b-values of 711 and 2855 s/mm^2 respectively. We choose to sample the higher b-value at twice the angular resolution of the lower b-value to account for its higher signal variation over the sphere due to its greater sensitivity to complex microstructure configurations. The optimized protocol thus consists of one shell with 30 gradient directions and b=711 s/mm^2 and the other with 60 directions and b=2855 s/mm^2. The protocol also includes 9 b=0 images.

2.3.2. Protocols for comparison

We acquire two additional HARDI shells with intermediate b-values and create alternative protocols by choosing from subsets of all the acquired shells. The two extra HARDI shells are a b=1000 s/mm^2 shell with 30 directions and a b=2000 s/mm^2 shell with 60 directions. All four HARDI shells share the same echo time (See Sec 2.4). All the protocols are tabulated in Table 1 and summarized below:

1. The *four-shell* protocol: Using all the available shells, this rich protocol provides pseudo ground-truth parameter estimates for quantifying the performance of the NODDI and alternative protocols.
2. The *two-shell* protocols: These protocols, one of which is the NODDI protocol (P14), are generated by choosing one 30-direction and one 60-direction shells. They differ from one another only in the choice of b-values.
3. The *reduced-orientation-sampling (ROS) NODDI* protocols: These are versions of the NODDI protocol requiring even shorter acquisition time. They are created by subdividing the orientations of each of its two shells uniformly into equal partitions. The optimal subdivision of the orientations is determined using the algorithm described in Cook et al. (2007) and implemented in Camino (Cook et al., 2006), which ensures the orientations within each subset are as uniformly distributed as possible.
4. The *one-shell* protocols: These represent the typical clinical protocols for diffusion tensor imaging.

2.4. Subject and data acquisition

2.4.1. Subject

Test data for *in vivo* imaging come from a healthy human volunteer (male, 35 years old). The subject was scanned with informed consent and the approval of the local research ethics committee.

2.4.2. Data acquisition

In vivo imaging is conducted on a clinical 3T Philips Achieva system with $|\mathbf{G}|_{max}$ = 65 mT/m. The HARDI shells are acquired using a pulsed-gradient spin-echo (PGSE) diffusion-weighted imaging (DWI) sequence with axial echo-planar imaging (EPI) readout. Different b-values are achieved by fixing the gradient pulse width δ to 17.5 ms and the pulse separation Δ to 37.8 ms while varying the gradient strength which takes the values 31.9, 37.8, 53.4, and 63.8 mT/m. The EPI readout uses a matrix size of 112×112 over a field of view (FOV) of 224×224 mm^2 and slice thickness of 2 mm, resulting in isotropic voxels of 2×2×2 mm^3. The echo time TE=78 ms and the repetition time TR=12.5 s are the same for all measurements. The signal to noise ratio (SNR) in white matter at b=0 is about 20. A total of 50 slices are acquired to cover the whole brain. The total scanning time for the NODDI protocol is about 25 minutes, with another 25 minutes for the two additional HARDI shells with intermediate b-values.

2.5. Synthetic data

To evaluate the NODDI and alternative protocols using data with known ground-truth tissue microstructure, we simulate MR signals from the orientation-dispersed white matter model described in Zhang et al. (2011). We choose this model because not only can it generate signals for tissues with different neurite density and orientation dispersion, but it also supports the prescription of realistic axon diameter values with its axon diameter parameter a. The latter enables us to assess the effect of approximating neurites as sticks.

Specifically, to simulate both gray and white matter, we set the true model parameter values to typical values for both tissue types. The set of values for each model parameter are tabulated in Table 2. We study each combination for a total of 80 different microstructural configurations. To account for potential orientation bias, we construct 250 instantiations of each configuration that differ from one another only in their mean orientations $\boldsymbol{\mu}$, which are sampled uniformly over the sphere. For each instantiation of each configuration, we synthesize the corresponding MR data using the 4-shell protocol and add synthetic Rician noise to match the SNR of the *in vivo* data.

2.6. Model fitting

2.6.1. Model parameters

The complete set of parameters for the NODDI model is

- ν_{ic}: intra-cellular volume fraction
- $d_{\|}$: intrinsic free diffusivity
- κ: concentration parameter of Watson distribution
- $\boldsymbol{\mu}$: mean orientation of Watson distribution
- ν_{iso}: isotropic volume fraction
- d_{iso}: isotropic diffusivity

As in Alexander et al. (2010) and Zhang et al. (2011), the diffusivities are fixed to their respective typical values *in vivo*: $d_{\|} = 1.7 \times 10^{-3}$ mm^{2}s^{-1} and $d_{iso} = 3.0 \times 10^{-3}$ mm^{2}s^{-1}. The remaining parameters are estimated from the fitting procedure described next. The estimated κ is then used to compute OD using Eq. (9).

2.6.2. Fitting routine

We fit the model to data with an adapted version of the routine described in Alexander et al. (2010). The procedure determines the maximum likelihood estimates of the parameters, using a Rician noise model, with the Gauss-Newton nonlinear optimization technique. The starting point for the nonlinear optimization is obtained from an initial brute-force search over a coarse and regular grid of physically plausible settings.

The only change we make is the omission of the additional MCMC procedure in Alexander et al. (2010). Although important for estimating weak parameters like the axon diameter index in (Alexander et al., 2010), for the NODDI model, the MCMC has negligible effect on the accuracy of our fitted model parameters. Its removal reduces the computation time dramatically, making parameter estimation over the whole-brain practical. On a standard workstation with two quad-core 3.0GHz Intel processors, by splitting the computation over 8 cores, the whole brain fitting requires no more than 3 hours.

2.7. Preprocessing

For the *in vivo* brain data, we manually delineate a binary mask defining the brain parenchyma using ITK-SNAP (Yushkevich et al., 2006). The brain mask defines the voxels for fitting with the routine described above.

To help identify overall trends in the microstructure parameter estimates for different brain tissues, we further segment the brain parenchyma into partitions of gray and white matter and CSF. We choose a simple segmentation procedure adequate for the purpose of assessing broad trends. Specifically, we define the CSF partition as the voxels with mean diffusivity, derived from diffusion tensor fit to the full imaging protocol, above 80% of d_{iso}, the expected value for CSF. For the remaining voxels, the gray matter partition is defined as the voxels with linearity (Westin et al., 2002) below 0.2, also estimated from diffusion tensor fit. The rest is then classified as the white matter partition.

3. Experiments and results

This section describes the synthetic and *in vivo* human data experiments for evaluating the performance of NODDI. The experiments are designed to quantify its accuracy and precision in estimating microstructure parameters with both the optimized and alternative protocols.

3.1. Synthetic data experiment

3.1.1. Design

We fit the NODDI tissue model to the subsets of the four-shell synthetic data. The subsets are defined in Table 1 and represent the NODDI and alternative protocols. The accuracy and precision of the microstructure parameter estimates from each protocol are assessed against the known ground-truth.

3.1.2. Results for intra-cellular and isotropic volume fractions

Fig. 2 plots the statistics of the estimated intra-cellular volume fractions ν_{ic} for different protocols. The figure illustrates the first key result of this work, that although ν_{ic} can not be estimated with the protocols using a single shell, it can be estimated with the ones using just two shells. Observe that the single-shell protocols consistently estimate ν_{ic} poorly, with both large upward bias and high variance, for the entire range of tested values. In contrast, all the multi-shell protocols estimate the ν_{ic} accurately and precisely, with only a slight upward bias. As expected, the protocol using all four shells results in the smallest bias and variance. The NODDI protocol (P14) and the other two-shell protocols are the next best. Here, the NODDI protocol does not show any appreciable advantage over the alternatives but does not demonstrate any disadvantage either.

The more surprising result is that the ROS-NODDI protocols perform almost as well as the NODDI protocol itself, with only a modest increase in the bias and variance. In particular, the ROS-NODDI protocols T1-3 collect no more orientations than any of the single-shell protocols. But by distributing these measurements over two shells, they can consistently estimate ν_{ic} while the single-shell protocols can not. Putting these observations together, it is evident that multi-shell protocols hold an advantage over their single-shell alternatives and that this advantage can be attributed solely to the use of multiple shells, rather than to any differences in orientation sampling density or SNR.

A similar pattern emerges when examining the statistics of the ν_{iso} estimates as shown in Fig. 3. The single-shell protocols result in both large upward bias and high variance, except for the highest ν_{ic}. The multi-shell protocols on the other hand lead to much smaller upward bias and variance.

The biases in ν_{ic} and ν_{iso} both increase as ν_{ic} decreases. This is expected because smaller ν_{ic} leads to larger and less anisotropic extra-cellular compartment, making the CSF contamination more difficult to separate from the neurite compartments.

3.1.3. Results for orientation dispersion and mean orientation

Fig. 4 plots the statistics for the orientation dispersion index OD. The figure illustrates the second key result, that OD can be estimated with protocols containing just a single shell. Overall, all the protocols demonstrate a similar level of performance. The four-shell protocol outperforms the others as expected but not by a large margin. The single-shell protocols P3 and P4 compare favorably to the two-shell protocols, suggesting that it is not essential to sample multiple b-values for estimating OD accurately.

The bias and variance of OD estimates depend primarily on the ground-truth value of OD itself. Specifically, when OD is low, below 0.5 in particular, the bias and variance are both very small; when OD is high, equal to or above 0.5, the variance increases substantially. When OD is very large, equal to or about 0.8, there is also a significant increase in negative bias. The dependence of the bias on the ground-truth value of OD mirrors the noise-induced bias in the estimate of FA (Pierpaoli and Basser, 1996; Jones and Basser, 2004): For low FA, corresponding to high OD, the bias is negative and its magnitude increases as the ground-truth value of FA decreases. Just as in DTI, this is unlikely to be a serious practical limitation. As Fig. 1 illustrates, the orientation distributions corresponding to very large ODs are not very different from one another and the high variance in its estimate simply reflects the lack of difference.

Despite the overall similarity in the performance across the protocols, some differences can be seen. In particular, when OD is high, the relative performance of a protocol appears to depend primarily on the number of sampled orientations rather than on the number of shells. On one hand, the effect of orientation sampling can be seen most clearly by comparing the NODDI and ROS-NODDI protocols. The latter perform worse than the NODDI protocol itself, with the level of performance degradation proportional to the level of the orientation sampling reduction. On the other hand, the minimal dependence on the number of shells can be seen by comparing the ROS-NODDI protocols T1-3 to the single-shell protocols P1 and P2. These two groups of protocols have the same number of orientations and show similar level of performance, despite having different numbers of shells. When OD is low, the only notable observation is that P1 and P2 are the worst performers. Again, a comparison to T1-3 suggests that sampling higher b-values is important for OD.

Fig. 5 plots the statistics for the estimates of the mean orientation $\boldsymbol{\mu}$, using the 95% cone of uncertainty (Jones, 2003). Similar to the main finding for orientation dispersion, the figure shows that the mean orientation can also be estimated with just one shell. Overall, all the protocols perform to similar levels of accuracy and precision that depend primarily on the value of OD. When OD is low, the mean orientation can be estimated accurately. But as OD increases, it becomes more difficult to estimate. At the extreme, when OD is equal to 1, the mean orientation is no longer defined and as expected, the cone of uncertainty takes its extreme value of 90^o. This dependence on OD is unsurprising and corresponds to the finding by Jones (2003) that the cone of uncertainty increases as the anisotropy decreases. As for orientation dispersion, difference in performance across the protocols is influenced more by difference in orientation sampling than by b-value.

3.1.4. Dependency on axon diameters

Estimates for all the microstructural parameters have minimal dependency on the axon diameters. This is assessed by grouping the tissue configurations according to their axon diameters and computing for each group the mean and standard deviation of the estimation error in each NODDI parameter. Using intracellular volume fraction as an example, for different axon diameters, the largest difference in the mean estimation error is less than 0.005 and the standard deviations for different axon diameters are almost identical.

3.2. In vivo human brain data experiment

3.2.1. Design

In the *in vivo* human data experiment, we use the microstructure parameter estimates derived from the four-shell protocol as the pseudo ground-truth and assess the accuracy and precision of the corresponding estimates computed from both the NODDI and alternative protocols.

3.2.2. Whole-brain parameter maps

Figs. 6 and 7 provide the whole-brain maps of the microstructure features estimated using the NODDI tissue model with the four-shell protocol. They are shown alongside the FA map, derived from fitting the diffusion tensor model to the same data, to illustrate qualitatively the utility of microstructure features for disentangling the factors contributing to the variation in FA across different anatomical regions of the brain.

The parameter maps from the NODDI model exhibit a spatial pattern of tissue distribution consistent with the known brain anatomy. The map of ν_{ic} shows the expected pattern of neurite density (Jespersen et al., 2010). It is lower in gray matter than in white matter; it takes its highest values in the major white matter tracts, such as the corpus callosum and the internal capsules. The map of OD demonstrates a trend that is similarly expected. It is higher in gray matter than in white matter; it takes its lowest values in the corpus callosum. This is also true for the map of ν_{iso}, which takes its highest values for the expected CSF regions, such as the ventricles.

The visual comparison of the FA map to the ν_{ic} and OD maps suggests that FA is primarily sensitive to OD and less so to ν_{ic}. The sensitivity of FA to OD can be seen most clearly from examining the white matter. While ν_{ic} does not vary substantially within the white matter, OD and FA exhibit significant regional variations that are inversely correlated with each other. At one end of the spectrum, in the corpus callosum where the axons are the most coherently-oriented, the lowest values of OD correspond to the highest values of FA. At the other end, in parts of the centrum semiovale, such as the highlighted region on slice 30 in Fig. 7, where there are significant fanning and crossing of axons, the highest values of OD coincide with the lowest values of FA, whereas ν_{ic} remains high throughout.

Although less apparent, many regional FA variations are associated with variations in both OD and ν_{ic}. We highlight two examples of such areas in Fig 6 on slice 6 and 22 respectively. In both regions, changes in FA are accompanied with changes in both OD and ν_{ic}. This illustrates that the two main factors contributing to FA can be disentangled with NODDI, allowing them to be studied separately.

3.2.3. Parameter correlations

Figs. 8 and 9 demonstrate the relationship between FA and the NODDI microstructure parameters, ν_{ic} and OD, quantitatively using scatterplots. The scatterplots are computed separately for the gray and white matter regions defined in Sec. 2.7. For both tissue types, FA shows a strong negative correlation to OD, consistent with the visual assessement above. For the white matter, FA exhibits a weaker positive correlation to ν_{ic}, which is not apparent from visual inspection. This combination of results suggest that, although the orientation dispersion is the most significant factor in determining the FA value, the neurite density also plays a non-negligible role.

To visualize the specific dependency of FA on ν_{ic} and OD, in Fig. 10 we plot the relationship between OD and ν_{ic} among the voxels with similar FA values. The figure confirms that, for both tissue types, a particular value of FA in tissue can be achieved by different combinations of ν_{ic} and OD. Furthermore, it reveals a sensible positive correlation between the two microstructure parameters, i.e., two voxels can have the same FA value as long as the one with the larger value of ν_{ic} also has the larger OD. Moreover, a small change in OD needs to be compensated with a much larger change in ν_{ic}, which is consistent with the observed relative sensitivity of FA to ν_{ic} and OD.

3.2.4. Visual assessment of protocol performance

Figs. 11 to 13 illustrate the parameter estimates from different protocols using a representative axial slice, slice 26 in Fig. 7. The observations from visual inspection agree with the two key results from the synthetic data experiment. In particular, Figs. 11 and 12 are consistent with the result concerning the two volume fraction parameters, i.e., that they can be estimated using the two-shell protocols, including the ROS-NODDI protocols, but not using the single-shell protocols. The maps of ν_{ic} and ν_{iso} using the two-shell protocols show close resemblance to those using the four-shell protocol. In contrast, the maps of ν_{ic} using the single-shell protocols have little resemblance to the one using the four-shell protocol and do not differentiate the gray matter from the white matter. The single-shell protocols estimate ν_{iso} slightly better than ν_{ic}. The maps of ν_{iso} vary from being very noisy but not significantly biased for low b-values (P1 and P2) to being substantially biased for high b-values (P3 and P4).

Fig. 13 confirms the result concerning the orientation dispersion parameter, i.e., that it can be estimated using just a single shell. The OD maps using both the single-shell and two-shell protocols are of very similar quality to the one using the four-shell protocol.

3.2.5. Quantitative assessment of protocol performance

To quantitatively confirm the observations from the visual assessment above, Figs. 14 to 17 plot the bias statistics of each microstructure parameter estimates from each protocol with respect to the four-shell protocol. These statistics are computed separately for the gray and white matter regions defined in Sec. 2.7. This is to account for the observation from the synthetic data experiment that the accuracy and precision of orientation dispersion and $\boldsymbol{\mu}$ depend strongly on the value of OD, which is significantly different between the gray and white matter.

The error statistics in estimating ν_{ic} and ν_{iso}, shown in Figs. 14 and 15 show overall good agreement with the synthetic data experiment finding for these two parameters. It is evident that, while the two-shell protocols provide accurate estimates of these parameters, the single-shell protocols give poor estimates that have large positive bias and variance. The errors for gray matter are particularly pronounced, consistent with the synthetic data experiment finding that the single-shell protocols estimate these parameters most poorly when ν_{ic} is low, which is the case in the gray matter.

These statistics also reveal something different from the synthetic data experiment. While the synthetic data experiment finds no clear difference among the two-shell protocols, the result here demonstrates otherwise. In particular, it shows that the NODDI protocol (P14) results in the smallest bias and variance, followed closely by the ROS-NODDI protocols. Among the two-shell protocols, the one with the largest bias is P13 and the one with the largest variance is P24.

The error statistics in estimating OD and $\boldsymbol{\mu}$, shown in Figs. 16 and 17, also demonstrate good consistency with the synthetic data experiment finding for these two parameters. It is clear that all protocols estimate these parameters to a similar accuracy, especially for white matter, suggesting that the number of shells is not the key factor that influences their estimation. For gray matter, consistent with the observation from the synthetic data experiment, the much larger OD leads to OD estimates that have larger variance and $\boldsymbol{\mu}$ estimates that are more biased. As the case for the synthetic data experiment, the single-shell protocols and the ROS-NODDI protocol lead to slightly more biased estimates of OD and more biased estimates of $\boldsymbol{\mu}$.

4. Discussion

In summary, this paper proposes and demonstrates NODDI, a practical technique for imaging neurite density and orientation dispersion on standard clinical MR systems. The technique combines the experiment design optimization in Alexander (2008) with a simple compartment model to identify an acquisition protocol that is straightforward to implement and sufficiently economical for clinical applications. The NODDI protocol consists of just two HARDI shells, the b-values of which can be easily achieved on clinical systems. The standard version acquires a whole-brain scan with 2mm isotropic resolution in about 25 minutes. By reducing angular resolution, the acquisition time can be shortened to under 10 minutes. Although this leads to the expected reduction in the accuracy of estimating the dominant fiber orientation, it has minimal effect on the estimation accuracy of the other neurite morphology parameters.

The key ingredient for the experiment design optimization is the NODDI tissue model, which simplifies the model in Zhang et al. (2011) to focus on the parameters that we can estimate reliably with a practical acquisition protocol. There, an orientation-dispersed white matter model is proposed to simultaneously estimate the axon diameter index (Alexander et al., 2010) and the orientation dispersion of axons. Estimating axon diameter, however, has onerous demands for both acquisition and fitting (Assaf et al., 2008; Barazany et al., 2009; Alexander et al., 2010; Zhang et al., 2011). Even the economical protocol for *in vivo* axon diameter estimation (Alexander et al., 2010; Zhang et al., 2011), also developed using the experiment design optimization, requires an hour to acquire half a brain at half the spatial resolution of the NODDI protocol, making it challenging for routine clinical use. Nevertheless, the demonstration of orientation dispersion estimation in Zhang et al. (2011) motivates us to determine a much more economical protocol by modeling

the axons as sticks rather than cylinders, i.e., by abandoning the axon diameter parameter. The simplified model not only leads to a clinically feasible protocol but also requires just a fraction of the computation time for estimating the model parameters.

Using NODDI, we demonstrate for the first time that it is possible to image both neurite density and orientation dispersion over the whole brain in a live subject, on a clinical scanner and in a clinically feasible acquisition time. This enables the disentanglement of two major factors contributing to FA and their separate analysis. By directly estimating the volume fraction of CSF, we are able to minimize the confounding effect of CSF-contamination, which affects periventricular white matter structures, such as the corpus callosum and the fornix (Metzler-Baddeley et al., 2012). Furthermore, we illustrate the importance of distangling these two factors with their independent influence on FA (Fig. 10). We additionally show that FA is influenced more strongly by orientation dispersion than neurite density, which is in good agreement with Jespersen et al. (2012).

We evaluate the accuracy and precision of NODDI extensively, using both simulated and *in vivo* data experiments, with the optimized protocol as well as with the alternative two-shell and single-shell protocols. The results from both experiments consistently demonstrate two novel findings in their own right. First, the neurite density cannot be estimated with single-shell data, regardless of the choice of b-value, but can be estimated with just two shells using clinically feasible b-values. Moreover, the comparison between the NODDI protocol and several alternative two-shell protocols demonstrate that the method does not depend strongly on the precise choice of the two b-values. Hence the proposed framework can be readily applied to existing multi-shell HARDI data sets acquired for other techniques, such as the CHARMED protocol (Assaf and Basser, 2005; De Santis et al., 2011) and hybrid diffusion imaging (HYDI) (Wu and Alexander, 2007), although the ideal protocol, for future studies, has just two shells with moderate b-values. Although we use a system with unusually high gradient strength of 65 mT/m here, NODDI should achieve similar performance on scanners with more standard ($|\mathbf{G}_{max}| = 40$ mT/m). Reducing $|\mathbf{G}_{max}|$ in the optimization reported in Sec 2.3 produces an optimized NODDI protocol still with two shells: the 30-direction shell has b-value of about 700 s/mm^2 and the 60-direction shell has b-value about 2000 s/mm^2. The reduced $|\mathbf{G}_{max}|$ shifts the balance of the trade-off between b-value and SNR toward lower maximum b-value to minimize the increase in TE: TE increases to 85 ms, which reduces SNR by about 10%. Simulations suggest the difference in protocols has minimal effect on the NODDI parameter estimation (See supplementary material). Second, the neurite orientation dispersion can be estimated with just one shell for sufficiently high b-values. In particular, this shows that even the standard clinical acquisition, single shell with $b = 1{,}000$ s/mm^2, can provide good estimates of the orientation dispersion, but lacks sensitivity for the neurite density. An important practical implication is that models such as Kaden et al. (2007), Sotiropoulos et al. (2012), and ours can be used with existing single-shell data, but further analysis should only include the orientation dispersion parameter, as the neurite density parameter is unreliable without the second shell.

The *in vivo* data experiment uses the parameter estimates from the four-shell protocol as the ground-truth. The results from the synthetic data experiment justify this choice. CHARMED and HYDI include shells with much higher b-values. However, simulation experiments using a six-shell protocol including two higher b-values of 5000 s/mm^2 and 6000 s/mm^2 show reduced performance compared to the four-shell protocol with the same total number of measurements (See supplementary material). The loss of signal due to increased TE to accommodate the larger b-values counteracts the potential benefits.

NODDI offers an opportunity to extend the application of neurite morphology quantification from being confined within the realm of postmortem histology to becoming a part of routine clinical practice. Jespersen et al. (2010) demonstrate a strong correlation of neurite density with the intensity of myelin stain under light microscopy, indicating that neurite density may be a useful marker for demyelination disorders, such as multiple sclerosis. The weaker dependence of FA to neurite density further suggests that neurite density may be a more sensitive marker of pathology than FA and may highlight early signs of demylelination before FA can.

The orientation dispersion index of neurites has two broad applications. In white matter, the orientation dispersion index quantifies the bending and fanning of axons, which is useful for mapping brain connectivity (Kaden et al., 2007). In particular, it will help determine whether many voxels currently classified as having crossing fibers (Jeurissen et al., 2010) may in fact consist of orientation-dispersed fibers with only a

single dominant orientation. Mapping orientaton dispersion over time may shed new light into the process of brain development (Jespersen et al., 2012). In gray matter, the index quantifies the pattern of sprawling dendritic processes. This provides a more direct marker of gray matter complexity than, for example, the measure derived from model-based bootstrapping proposed in Haroon et al. (2010), which has the potential in differentiating neurodegenerative disorders of varying severity (Haroon et al., 2011).

NODDI can be improved in a number of ways. First, in the current demonstration, the NODDI tissue model currently parametrizes the neurite orientation distribution with Watson distribution. This has limited accuracy in modeling orientation distributions that are not cylindrically symmetric, such as in the regions with fanning or crossing axons. Although beyond the scope of the present work, NODDI extends naturally to model these more complex orientation distributions using, e.g., the Bingham distribution (Mardia and Jupp, 1990) for fanning configurations and mixtures of Watson or Bingham distributions for crossing configurations. Earlier work (Kaden et al., 2007; Sotiropoulos et al., 2012) demonstrates the feasibility of such an extension. The resulting model selection problem can be solved using a variety of existing strategies with many successful examples in the diffusion MRI literature (Alexander et al., 2002; Hosey et al., 2005; Behrens et al., 2007; Kaden et al., 2007; De Santis et al., 2011; Sotiropoulos et al., 2012). Second, the NODDI protocol is currently derived without explicitly optimizing for the orientation sampling scheme across multiple HARDI shells. Minor improvements may come from multi-shell orientation sampling optimization approaches, such as those recently proposed in Caruyer et al. (2011); De Santis et al. (2011).

Acknowledgment

We would like to thank our anonymous reviewers whose comments are invaluable and have led to significant improvement of this manuscript. This work is supported by the future and emerging technologies (FET) program of the EU FP7 framework through the CONNECT consortium (www.brain-connect.eu), and the MS Society of Great Britain and Northern Ireland, the ISRT and the CBRC. DCA is additionally funded by EPSRC under grant EP/E007748.

References

Abramowitz, M., Stegun, I. A., 1972. Handbook of mathematical functions with formulas, graphs, and mathematical tables, 9th Edition. New York: Dover.

Alexander, D. C., 2008. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. Magn. Reson. Med. 60, 439–448.

Alexander, D. C., Barker, G. J., Arridge, S. R., 2002. Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. Magn. Reson. Med. 48, 331–340.

Alexander, D. C., Hubbard, P. L., Hall, M. G., Moore, E. A., Ptito, M., Parker, G. J. M., Dyrby, T. B., Oct 2010. Orientationally invariant indices of axon diameter and density from diffusion MRI. NeuroImage 52 (4), 1374–1389.

Assaf, Y., Basser, P. J., 2005. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. NeuroImage 27, 48–58.

Assaf, Y., Blumenfeld-Katzir, T., Yovel, Y., Basser, P. J., 2008. AxCaliber: a method for measuring axon diameter distribution from diffusion MRI. Magn. Reson. Med. 59, 1347–1354.

Assaf, Y., Cohen, Y., 2000. Assignment of the water slow-diffusing component in the central nervous system using q-space diffusion MRS: implications for fiber tract imaging. Magn. Reson. Med. 43, 191–199.

Assaf, Y., Cohen, Y., 2009. Inferring microstructural information of white matter from diffusion MRI. In: Johansen-Berg, H., Behrens, T. E. J. (Eds.), Diffusion MRI: From Quantitative Measurement to In-vivo Neuroanatomy. Academic Press, pp. 127–146.

Assaf, Y., Freidlin, R. Z., Rhode, G. K., Basser, P. J., 2004. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. Magn. Reson. Med. 52, 965–978.

Barazany, D., Basser, P. J., Assaf, Y., 2009. In-vivo measurement of the axon diameter distribution in the corpus callosum of a rat brain. Brain 132, 1210–1220.

Basser, P. J., Mattiello, J., Bihan, D. L., 1994. MR diffusion tensor specstroscopy and imaging. Biophys. J. 66, 259–267.

Beaulieu, C., 2009. The biological basis of diffusion anisotropy. In: Johansen-Berg, H., Behrens, T. E. J. (Eds.), Diffusion MRI: From Quantitative Measurement to In-vivo Neuroanatomy. Academic Press, pp. 105–126.

Behrens, T. E. J., Johansen-Berg, H., Jbabdi, S., Rushworth, M. F. S., Woolrich, M. W., 2007. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? NeuroImage 34, 144–155.

Behrens, T. E. J., Johansen-Berg, H., Woolrich, M. W., Smith, S. M., Wheeler-Kingshott, C. A. M., Boulby, P. A., Barker, G. J., Sillery, E. L., Sheehan, K., Cicarelli, O., Thompson, A. J., Brady, J. M., Matthews, P. M., 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. Magn. Reson. Med. 50, 1077–1088.

Bodini, B., Ciccarelli, O., 2009. Diffusion MRI in neurological disorders. In: Johansen-Berg, H., Behrens, T. E. J. (Eds.), Diffusion MRI: From Quantitative Measurement to In-vivo Neuroanatomy. Academic Press, pp. 175–204.

Bruijn, L. I., Miller, T. M., Cleveland, D. W., 2004. Unraveling the mechanisms involved in motor neuron degeneration in ALS. Annual Review of Neuroscience 27, 723–749.

Bürgel, U., Amunts, K., Hoemke, L., Mohlberg, H., Gilsbach, J. M., Zilles, K., 2006. White matter fiber tracts of the human brain: Three-dimensional mapping at microscopic resolution, topography and intersubject variability. NeuroImage 29, 1092–1105.

Caruyer, E., Lenglet, C., Sapiro, G., Deriche, R., 2011. Incremental gradient table for multiple Q-shells diffusion MRI. In: Proceedings Human Brain Mapping.

Conel, J. L., 1939. The postnatal development of the human cerebral cortex. Harvard University Press, Cambridge, USA.

Cook, P. A., Bai, Y., Nedjati-Gilani, S., Seunarine, K. K., Hall, M. G., Parker, G. J., Alexander, D. C., May 2006. Camino: Open-source diffusion-MRI reconstruction and processing. In: 14th Scientific Meeting of the Internationaal Society for Magnetic Resonance in Medicine. p. 2759.

Cook, P. A., Symms, M., Boulby, P. A., Alexander, D. C., 2007. Optimal acquisition orders of diffusion-weighted MRI measurement. J. Magn. Reson. Imaging 25 (5), 1051–1058.

De Santis, S., assaf, Y., Evans, C. J., Jones, D. K., 2011. Improved precision in the charmed model of white matter through sampling scheme optimization and model parsimony testing. In: Proceedings Intl Soc Magn Res Med. pp. 3928–3928.

Evanglou, N., Esiri, M. M., Smith, S., Palace, J., Matthews, P. M., 2000. Quantitative pathological evidence for axonal loss in normal appearing white matter in multiple sclerosis. Ann Neurol 47, 391–395.

Haroon, H. A., Binney, R. J., Parker, G. J., 2010. Probabilistic quantification of regional cortical microstructural complexity. In: Proceedings Intl Soc Magn Res Med. pp. 578–578.

Haroon, H. A., Reynolds, H., Carter, S. F., Embleton, K. V., Parker, K. G. H. G. J., 2011. HARDI-based microstructural complexity mapping reveals distinct subcortical and cortical grey matter changes in mild cognitive impairment and alzheimer's disease. In: Proceedings Intl Soc Magn Res Med. pp. 682–682.

Hosey, T., Williams, G., Ansorge, R., 2005. Inference of multiple fiber orientations in high angular resolution diffusion imaging. Magn. Reson. Med. 54, 1480–1489.

House, E. L., Pansky, B., 1960. A functional approach to neuroanatomy. New York: McGraw-Hill.

Jacobs, B., Driscoll, L., Schall, M., 1997. Life-span dendritic and spine changes in areas 10 and 18 of human cortex: a quantitative Golgi study. J Comparative Neurology 386, 661–680.

Jacobs, B., Schall, M., Prather, M., Kapler, E., Driscoll, L., Baca, S., Jacobs, J., Ford, K., Wainwright, M., Treml, M., Jun 2001. Regional dendritic and spine variation in human cerebral cortex: a quantitative Golgi study. Cerebral Cortex 11, 558–571.

Jespersen, S. N., Bjarkam, C. R., Nyengaard, J. R., Chakravarty, M. M., Hansen, B., Vosegaard, T., Ostergaard, L., Yablonskiy, D. A., Chr. Nielsen, N., Vestergaard-Poulsen, P., 2010. Neurite density from magnetic resonance diffusion measurements at ultrahigh field: comparison with light microscopy and electron microscopy. NeuroImage 49, 205–216.

Jespersen, S. N., Kroenke, C. D., Ostergaard, L., Ackerman, J. J. H., Yablonskiy, D. A., 2007. Modeling dendrite density from magnetic resonance diffusion measurements. NeuroImage 34, 1473–1486.

Jespersen, S. N., Leigland, L. A., Cornea, A., Kroenke, C. D., 2012. Determination of axonal and dendritic orientation distributions within the developing cerebral cortex by diffusion tensor imaging. IEEE Trans. Med. Imaging 31 (1), 16–32.

Jeurissen, B., Leemans, A., Tournier, J.-D., Sijbers, J., 2010. Estimating the number of fiber orientations in diffusion MR voxels: a constrained spherical deconvolution study. In: Proc. ISMRM.

Jones, D. K., 2003. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. Magn. Reson. Med. 49, 7–12.

Jones, D. K., Basser, P. J., 2004. "Squashing peanuts and smashing pumpkins": how noise distorts diffusion-weighted MR data. Magn. Reson. Med. 52, 979–993.

Kaden, E., Knösche, T. R., Anwander, A., 2007. Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion MR imaging. NeuroImage 37, 474–488.

Le Bihan, D. (Ed.), 1995. Diffusion and perfusion magnetic resonance imaging: applications to functional MRI. Raven Press, New York.

Mardia, K. V., Jupp, P. E., 1990. Directional statistics. Wiley series in probability and statistics. John Wiley & Sons, Ltd.

Metzler-Baddeley, C., O'Sullivan, M. J., Bells, S., Pasternak, O., Jones, D. K., 2012. How and how not to correct for CSF-contamination in diffusion MRI. NeuroImage 59 (2), 1394–1403.

Panagiotaki, E., Schneider, T., Siow, B., Hall, M. G., Lythgoe, M. F., Alexander, D. C., 2012. Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison. NeuroImage 59 (3), 2241–2254.

Paula-Barbosa, M. M., Mota Cardoso, R., Guimaraes, M. L., Cruz, C., 1980. Dendritic degeneration and regrowth in the cerebral cortex of patients with Alzheimer's disease. J Neurological Sciences 45 (1), 129–134.

Pierpaoli, C., Basser, P. J., 1996. Toward a quantitative assessment of diffusion anisotropy. Magn. Reson. Med. 36, 893–906.

Pierpaoli, C., Jezzard, P., Basser, P. J., Barnett, A., Chiro, G. D., 1996. Diffusion tensor MR imaging of the human brain. Radiology 201, 637–648.

Salat, D. H., Lee, S. Y., Yu, P., Setty, B., Rosas, H. D., Grant, P. E., 2009. DTI in development and aging. In: Johansen-Berg, H., Behrens, T. E. J. (Eds.), Diffusion MRI: From Quantitative Measurement to In-vivo Neuroanatomy. Academic Press, pp. 205–236.

Sotiropoulos, S. N., Behrens, T. E., Jbabdi, S., 2012. Ball and rackets: Inferring fiber fanning from diffusion-weight MRI. NeuroImage 60 (2), 1412–1425.

Stanisz, G. J., Szafer, A., Wright, G. A., Henkelman, M., 1997. An analytical model of restricted diffusion in bovine optic

nerve. Magn. Reson. Med. 37, 103–111.
Szafer, A., Zhong, J. H., Gore, J. C., 1995. Theoretical model for water diffusion in tissues. Magn. Reson. Med. 33 (5), 697–712.
Westin, C.-F., Maier, S. E., Mamata, H., Nabavi, A., Jolesz, F. A., Kikinis, R., 2002. Processing and visualization for diffusion tensor MRI. Medical Image Analysis 6, 93–108.
Wu, Y. C., Alexander, A. L., July 2007. Hybrid diffusion imaging. NeuroImage 36 (3), 617–629.
Yushkevich, P. A., Piven, J., Hazlett, H. C., Smith, R. G., Ho, S., Gee, J. C., Gerig, G., 2006. User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. NeuroImage 31 (3), 1116–1128.
Zhang, H., Hubbard, P. L., Parker, G. J., Alexander, D. C., June 2011. Axon diameter mapping in the presence of orientation dispersion with diffusion MRI. NeuroImage 56 (3), 1301–1315.

Tables

Protocol	Settings
Pall/All	the full data set
P1	b=711 (30)
P2	b=1000 (30)
P3	b=2000 (60)
P4	b=2855 (60)
P13	b=711 (30) & b=2000 (60)
P14	b=711 (30) & b=2855 (60)
P23	b=1000 (30) & b=2000 (60)
P24	b=1000 (30) & b=2855 (60)
S1/S2	b=711 (15) & b=2855 (30)
T1/T2/T3	b=711 (10) & b=2855 (20)

Table 1: The list of imaging protocols used for evaluation. The b-values are expressed in the unit of s/mm^2. The number of sampled orientations for a particular b-value is shown in parentheses. The NODDI protocol (P14) and its reduced-orientation-sampling (ROS) versions (S1-2, T1-3) are highlighted. S1 and S2 denote the ROS subsets from splitting the NODDI protocol, into two halves. Similarly, T1, T2, and T3 denote the ROS subsets from splitting the NODDI protocol into three equal partitions.

Parameter	Ground-truth values
ν_{ic}	$\{0.2, 0.4, 0.6, 0.8\}$
ν_{iso}	$\{0.0\}$
a	$\{0.5, 1, 2, 4\}\mu$m
κ	$\{0, 0.25, 1, 4, 16\}$
$\boldsymbol{\mu}$	250 uniformly distributed orientations

Table 2: Ground-truth parameter values for the synthetic tissue substrates.

Figure Captions

Figure 1. Illustration of a set of Watson distributions with the same mean orientation but different orientation dispersion index: OD $\in \{0.04, 0.16, 0.5, 0.84, 1.0\}$. The Watson distribution is cylindrically symmetric, hence showing only the cross-sectional view through the symmetry axis which coincides with its mean orientation. The sampled OD values are chosen to correspond to the sampled concentration parameters used for generating the synthetic data set as given in Table 2.

Figure 2. Statistics of the estimates of the intra-cellular volume fraction ν_{ic} from the synthetic data experiment using different protocols. For each of the four ground-truth ν_{ic} values, the statistics are pooled over the tissue configurations with that ν_{ic} value and shown in its respective panel. In its panel, we plot, for each protocol, the mean and standard deviation of ν_{ic} estimates derived from the corresponding set of tissue configurations that vary in axon diameters, orientation dispersions, as well as mean orientations. The ground-truth value is indicated with a dashed line. The key to the protocols is in Table 1.

Figure 3. As Fig. 2 but plotting the statistics of the estimates of the isotropic (CSF) volume fraction ν_{iso}. Different panels show the statistics of the subsets of the tissue configurations with identical ν_{ic} as defined in Fig. 2. Because the ground-truth value of ν_{iso} is 0 for all the substrates, the dashed line used to indicate this value is obscured by the x-axis.

Figure 4. As Fig. 2 but plotting the statistics of the estimates of orientation dispersion index OD. Different from Fig. 2, the tissue configurations are grouped according to their values of OD, which are specified via the concentration parameter κ.

Figure 5. As Fig. 4 but plotting the statistics of the estimates of mean orientation $\boldsymbol{\mu}$, using the 95% cone of uncertainty (Jones, 2003). The ground-truth cone of uncertainty is zero.

Figure 6. Maps of RGB-encoded principal direction $\boldsymbol{\mu}$, FA, orientation dispersion index OD, intra-cellular volume fraction ν_{ic}, and isotropic (CSF) volume fraction ν_{iso} computed using the four-shell data, showing every 4th slice of the inferior half of the brain. Please refer to the text for the explanation of the regions highlighted with dashed ellipses.

Figure 7. As Fig. 6 but for the superior half of the brain.

Figure 8. Scatterplots of FA vs ν_{ic} for gray (left) and white (right) matters.

Figure 9. Scatterplots of FA vs OD for gray (left) and white (right) matters.

Figure 10. Scatterplots of ν_{ic} vs OD for gray (left) and white (right) matter. To assess the relationship between ν_{ic} and OD for the voxels with similar FA values, we stratify the voxels into smaller groups each of which has their FA values falling between a narrow range specified in the legend.

Figure 11. Maps of intra-cellular volume fraction ν_{ic} from different protocols for visual comparison. For the ROS-NODDI protocols, only one example from each protocol is shown.

Figure 12. As Fig. 11 but showing maps of isotropic (CSF) volume fraction ν_{iso}.

Figure 13. As Fig. 11 but showing maps of orientation dispersion index OD.

Figure 14. Mean and standard deviation of the estimation errors in ν_{ic} using various protocols when benchmarked against the estimates using the full protocol. The error statistics are calculated separately for gray matter (left) and white matter (right).

Figure 15. As Fig. 14 but plotting the error statistics of ν_{iso}.

Figure 16. As Fig. 14 but plotting the error statistics of orientation dispersion.

Figure 17. As Fig. 14 but plotting the error statistics of $\boldsymbol{\mu}$.

Figures

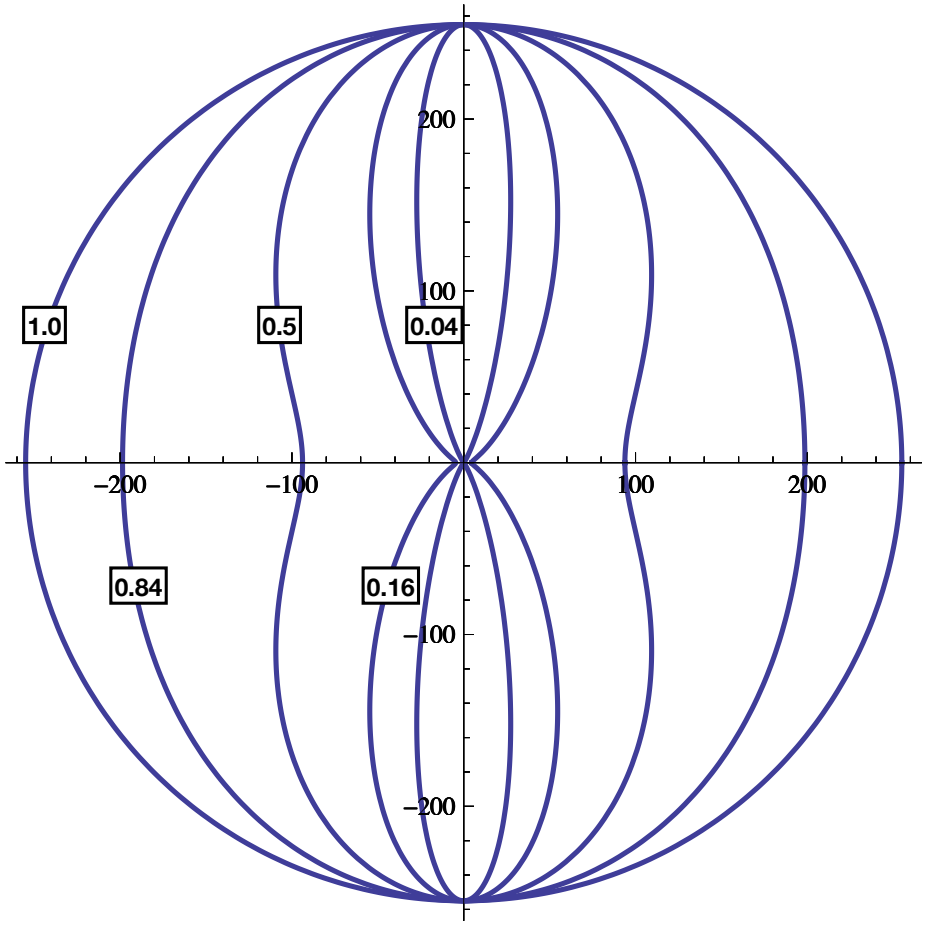

Figure 1: Illustration of a set of Watson distributions with the same mean orientation but different orientation dispersion index: OD $\in \{0.04, 0.16, 0.5, 0.84, 1.0\}$. The Watson distribution is cylindrically symmetric, hence showing only the cross-sectional view through the symmetry axis which coincides with its mean orientation. The sampled OD values are chosen to correspond to the sampled concentration parameters used for generating the synthetic data set as given in Table 2.

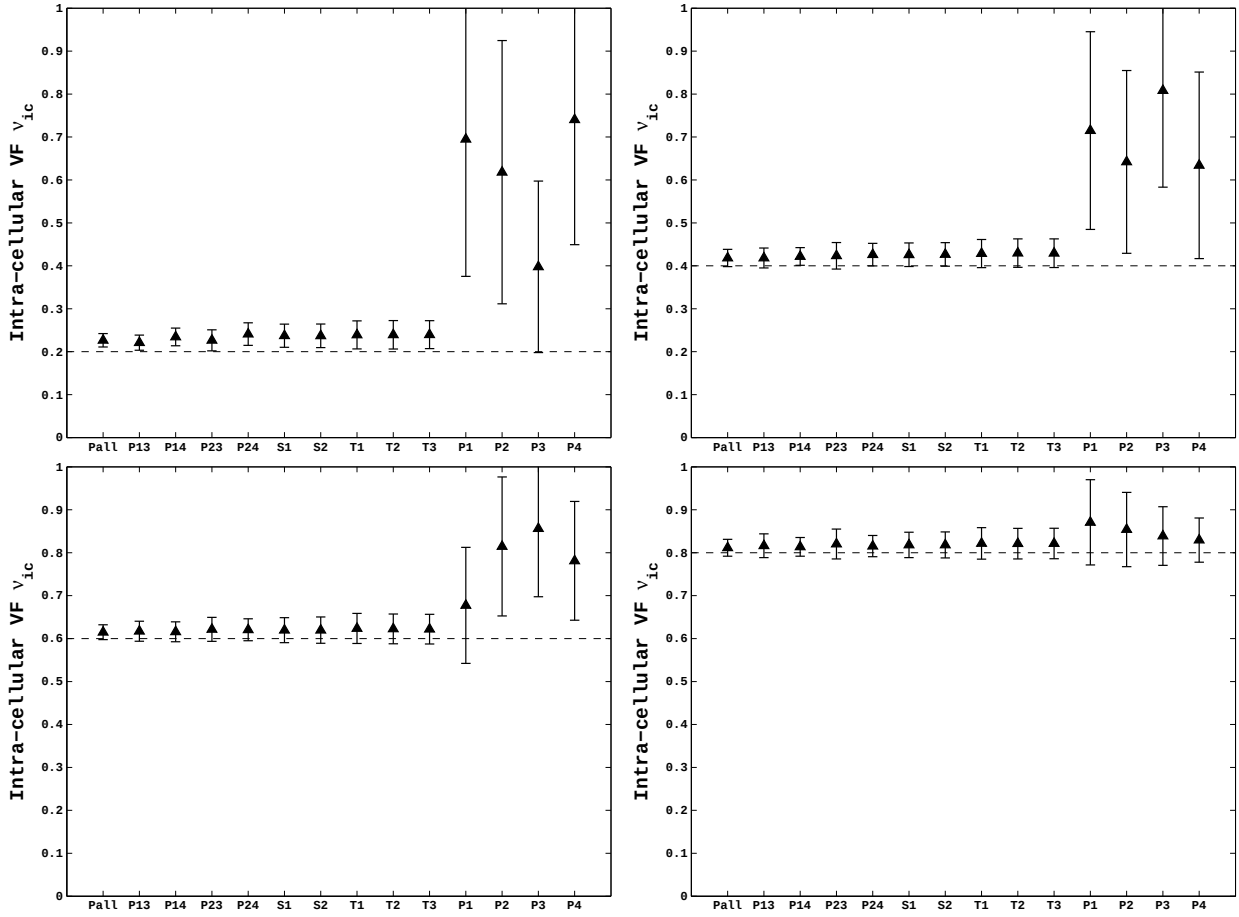

Figure 2: Statistics of the estimates of the intra-cellular volume fraction ν_{ic} from the synthetic data experiment using different protocols. For each of the four ground-truth ν_{ic} values, the statistics are pooled over the tissue configurations with that ν_{ic} value and shown in its respective panel. In its panel, we plot, for each protocol, the mean and standard deviation of ν_{ic} estimates derived from the corresponding set of tissue configurations that vary in axon diameters, orientation dispersions, as well as mean orientations. The ground-truth value is indicated with a dashed line. The key to the protocols is in Table 1.

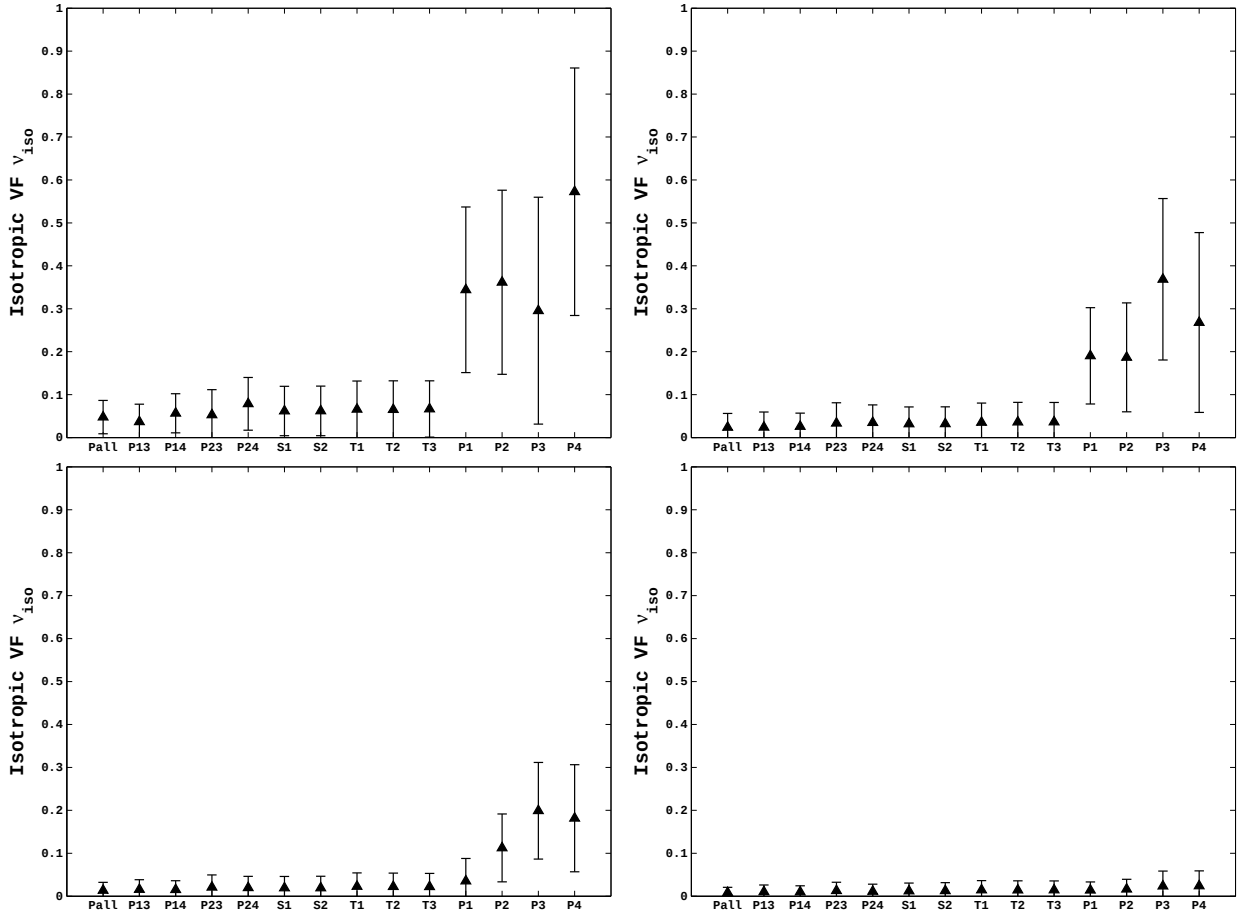

Figure 3: As Fig. 2 but plotting the statistics of the estimates of the isotropic (CSF) volume fraction ν_{iso}. Different panels show the statistics of the subsets of the tissue configurations with identical ν_{ic} as defined in Fig. 2. Because the ground-truth value of ν_{iso} is 0 for all the substrates, the dashed line used to indicate this value is obscured by the x-axis.

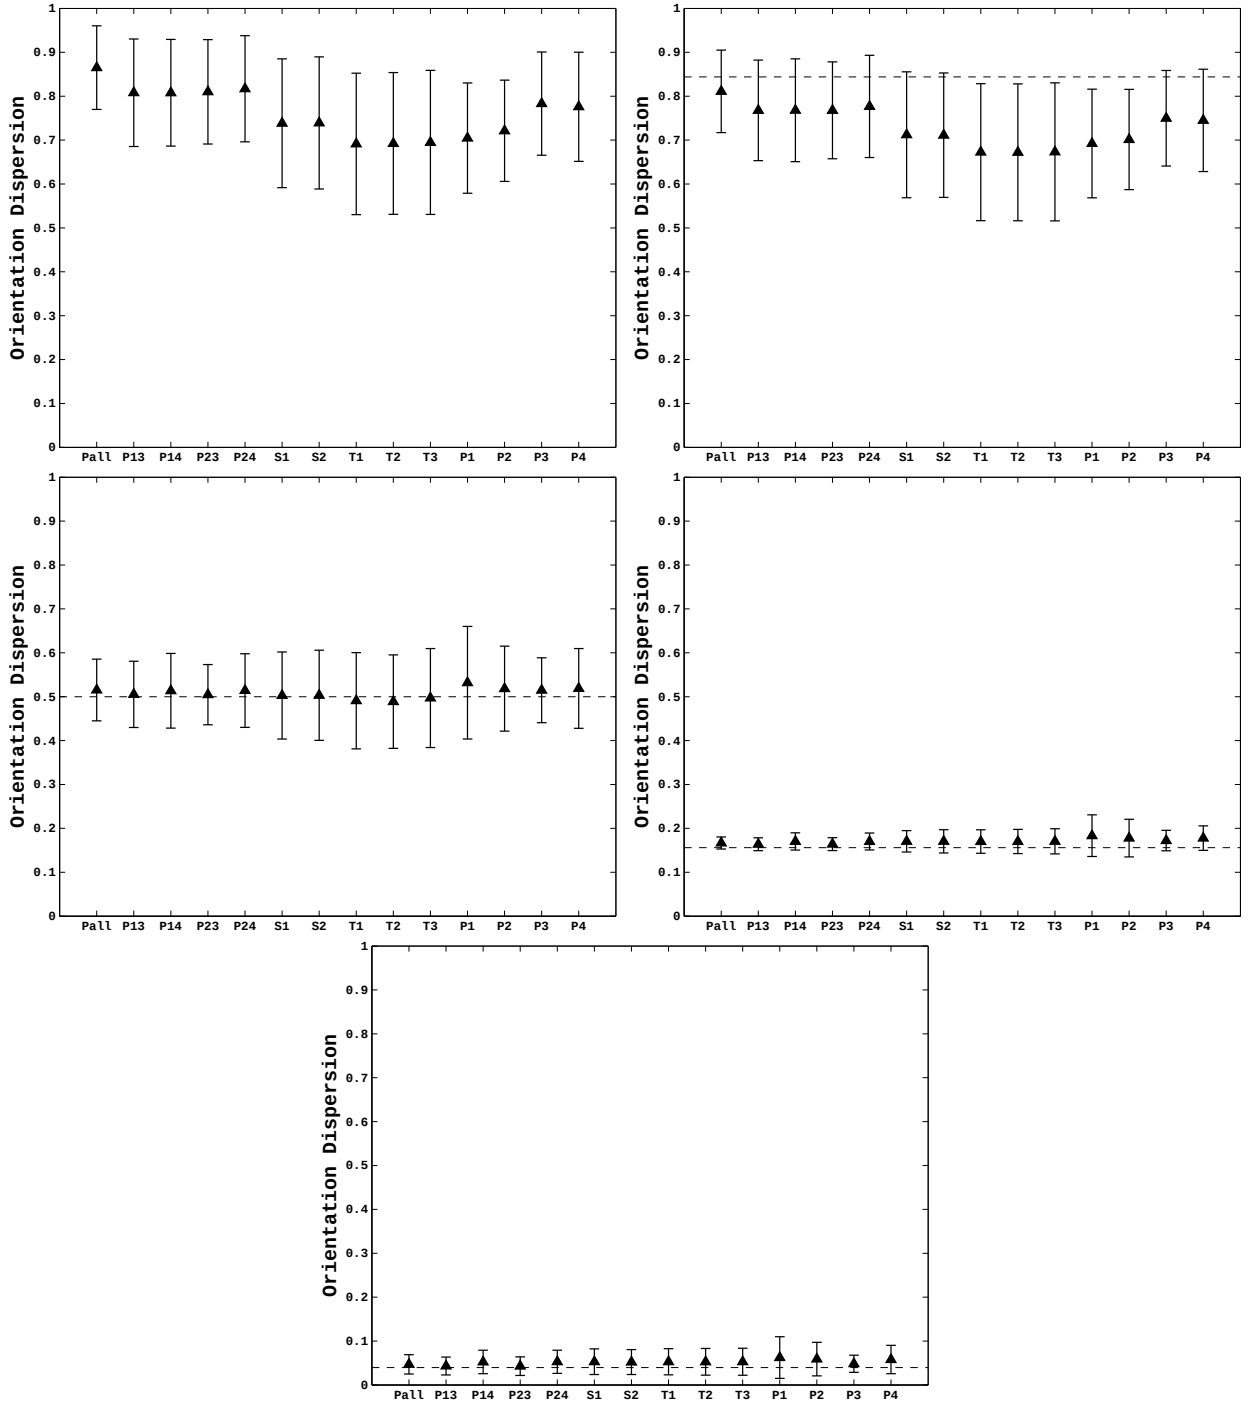

Figure 4: As Fig. 2 but plotting the statistics of the estimates of orientation dispersion index OD. Different from Fig. 2, the tissue configurations are grouped according to their values of OD, which are specified via the concentration parameter κ.

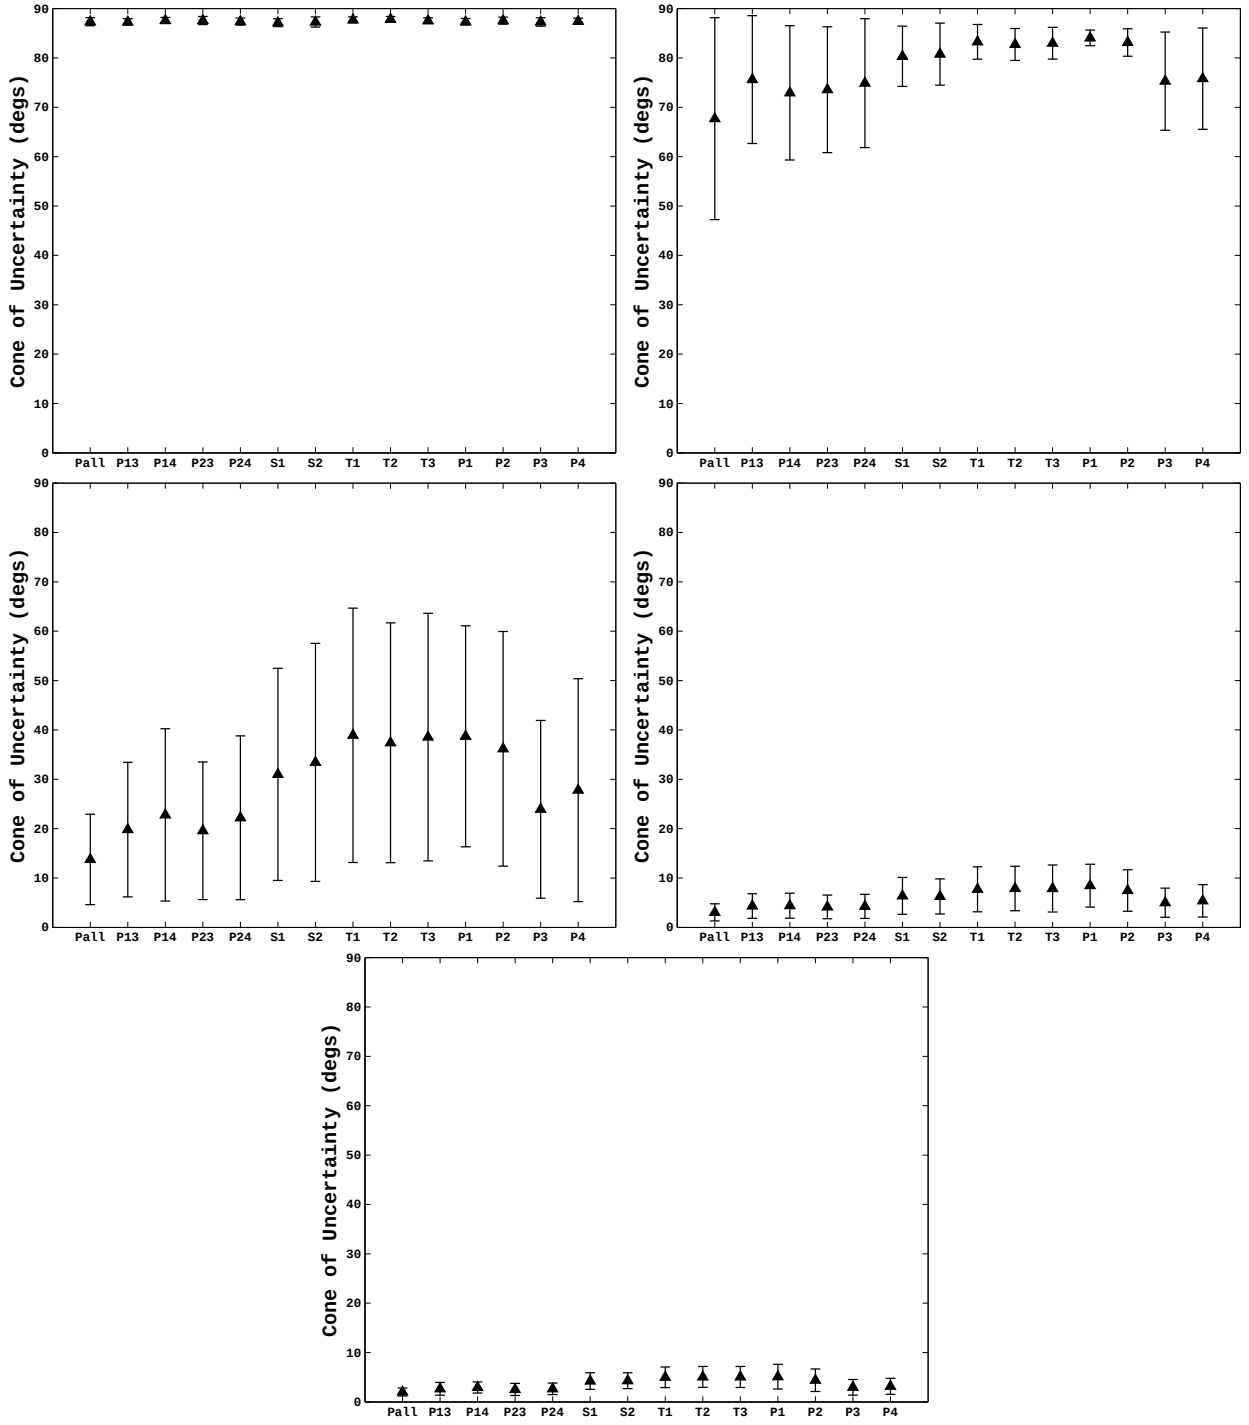

Figure 5: As Fig. 4 but plotting the statistics of the estimates of mean orientation $\boldsymbol{\mu}$, using the 95% cone of uncertainty (Jones, 2003). The ground-truth cone of uncertainty is zero.

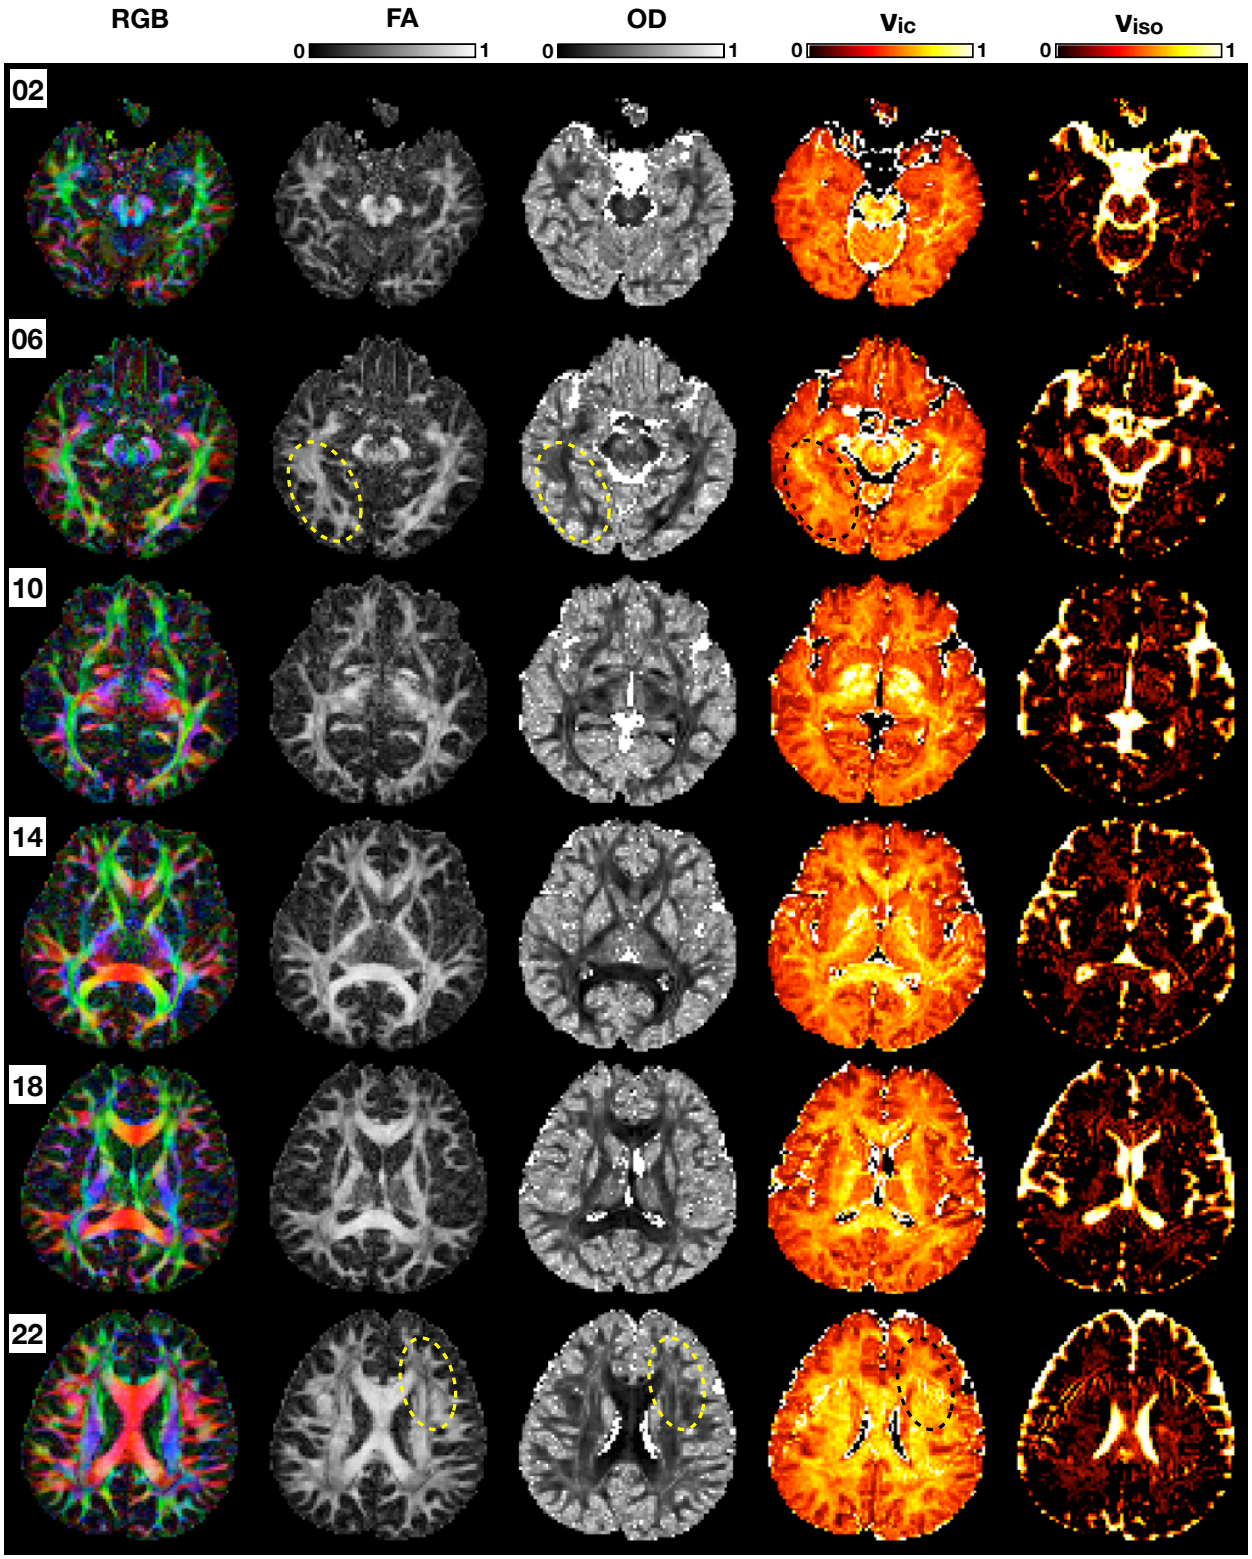

Figure 6: Maps of RGB-encoded principal direction $\boldsymbol{\mu}$, FA, orientation dispersion index OD, intra-cellular volume fraction ν_{ic}, and isotropic (CSF) volume fraction ν_{iso} computed using the four-shell data, showing every 4th slice of the inferior half of the brain. Please refer to the text for the explanation of the regions highlighted with dashed ellipses.

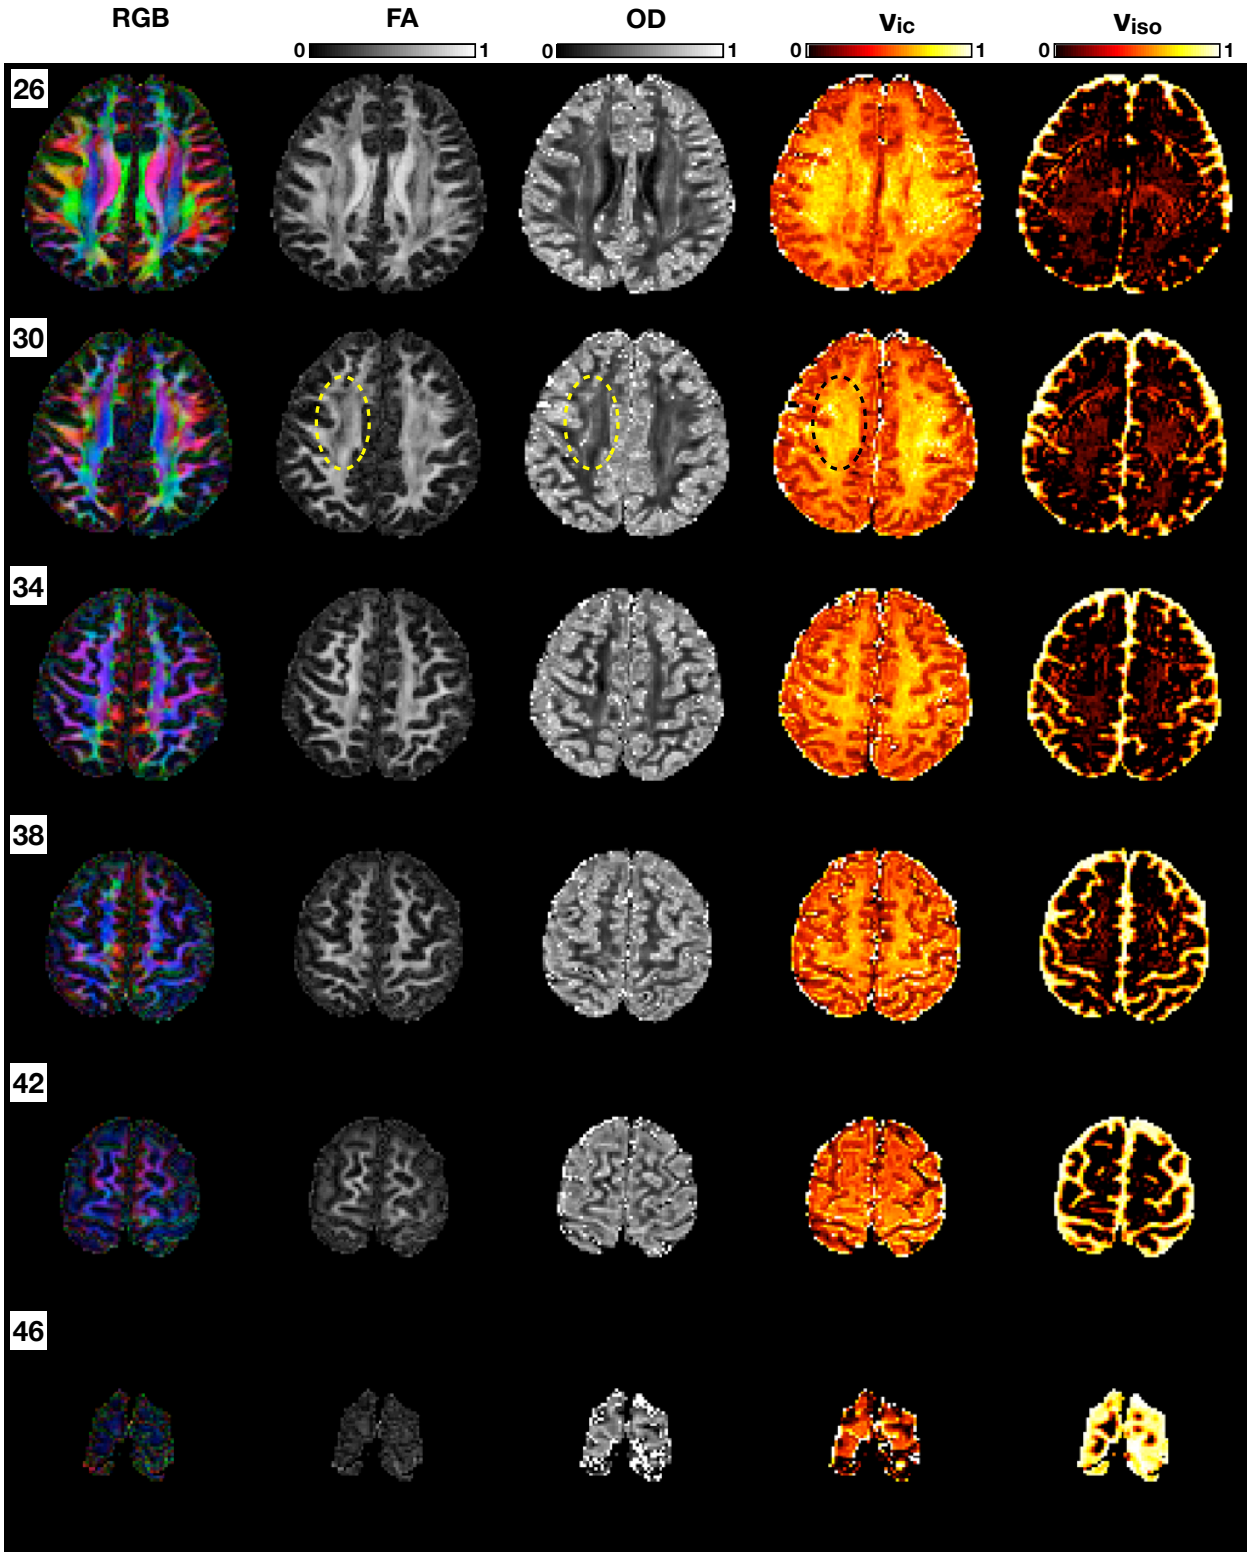

Figure 7: As Fig. 6 but for the superior half of the brain.

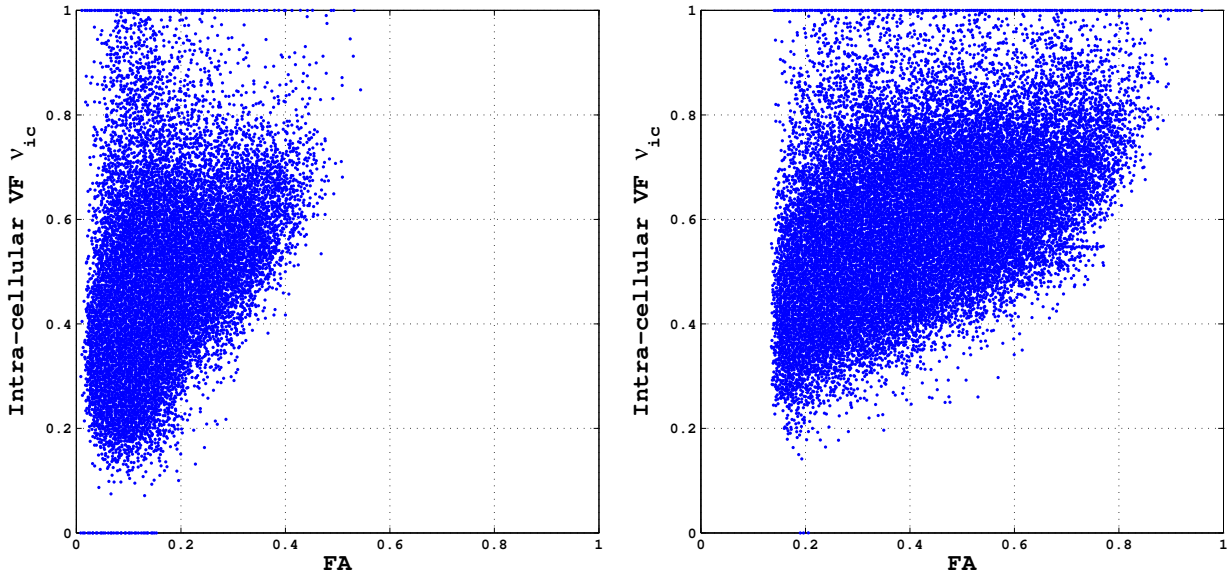

Figure 8: Scatterplots of FA vs ν_{ic} for gray (left) and white (right) matters.

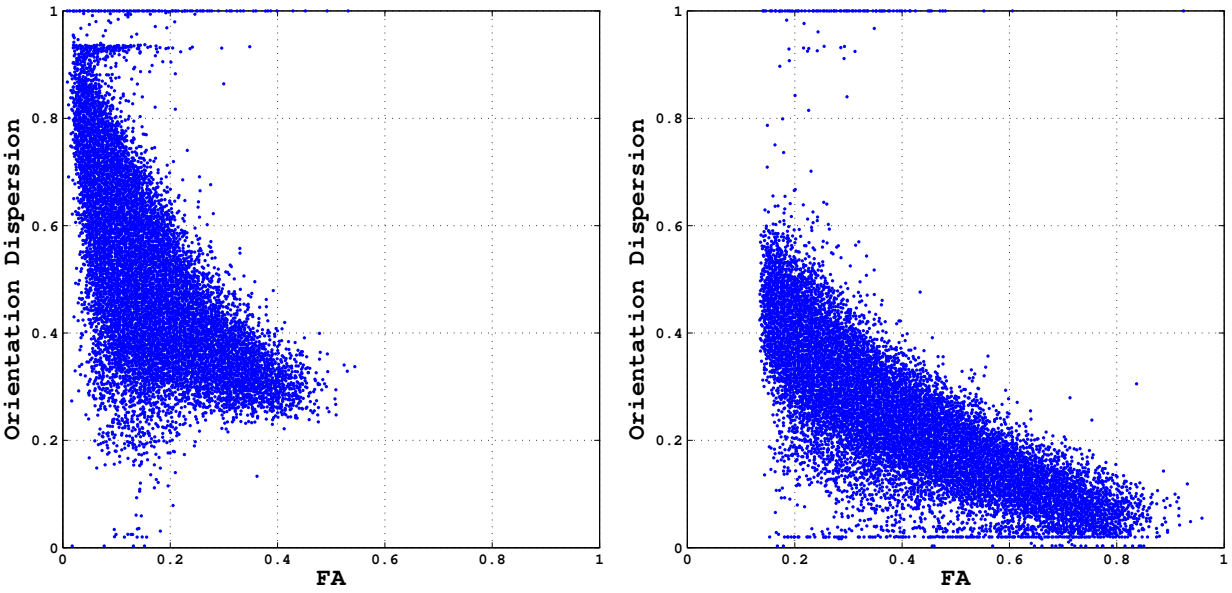

Figure 9: Scatterplots of FA vs OD for gray (left) and white (right) matters.

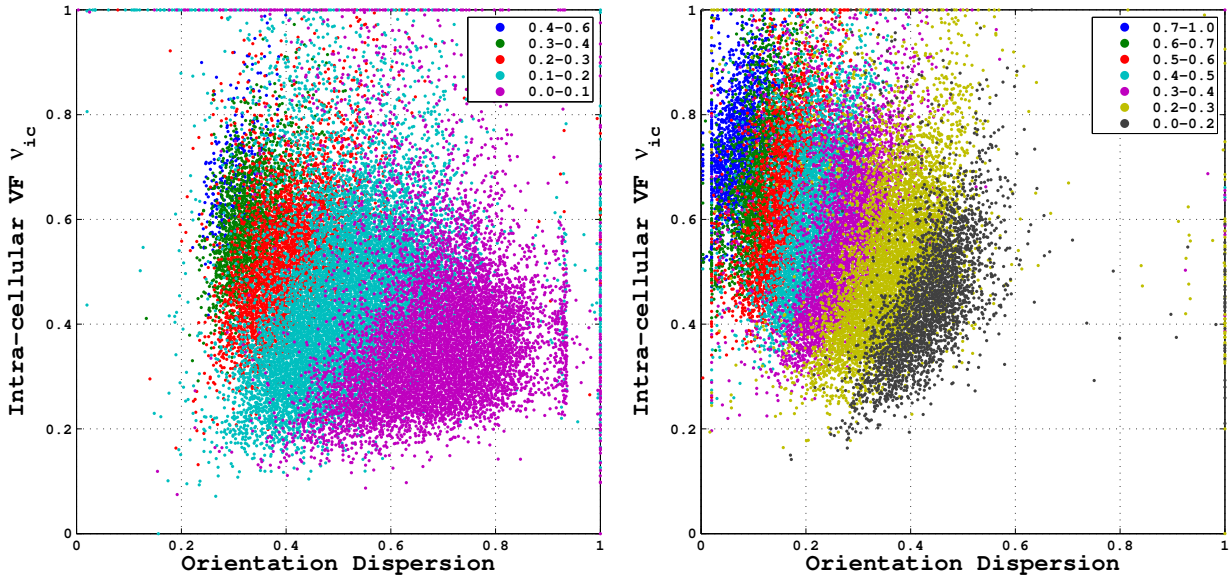

Figure 10: Scatterplots of ν_{ic} vs OD for gray (left) and white (right) matter. To assess the relationship between ν_{ic} and OD for the voxels with similar FA values, we stratify the voxels into smaller groups each of which has their FA values falling between a narrow range specified in the legend.

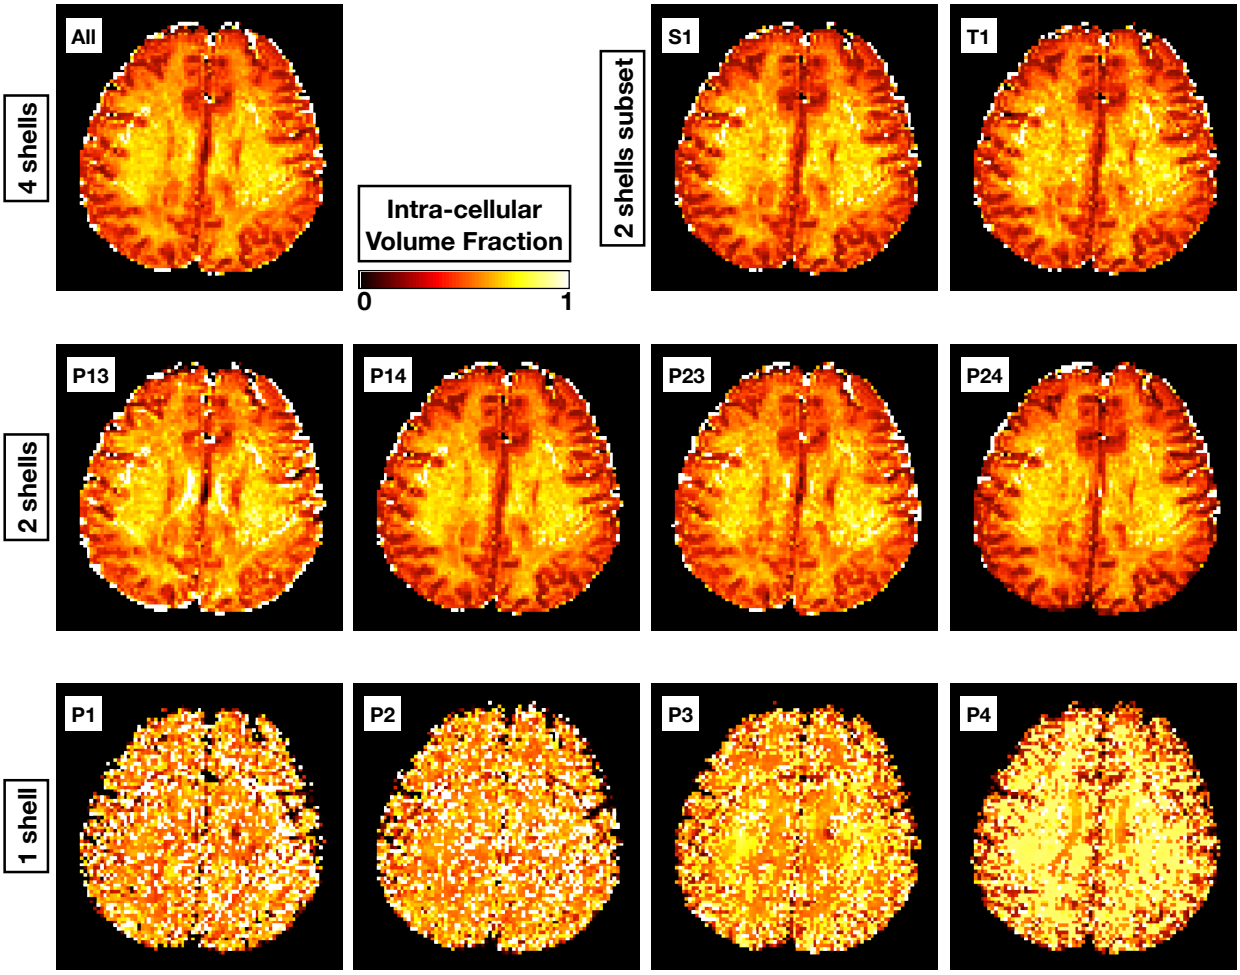

Figure 11: Maps of intra-cellular volume fraction ν_{ic} from different protocols for visual comparison. For the ROS-NODDI protocols, only one example from each protocol is shown.

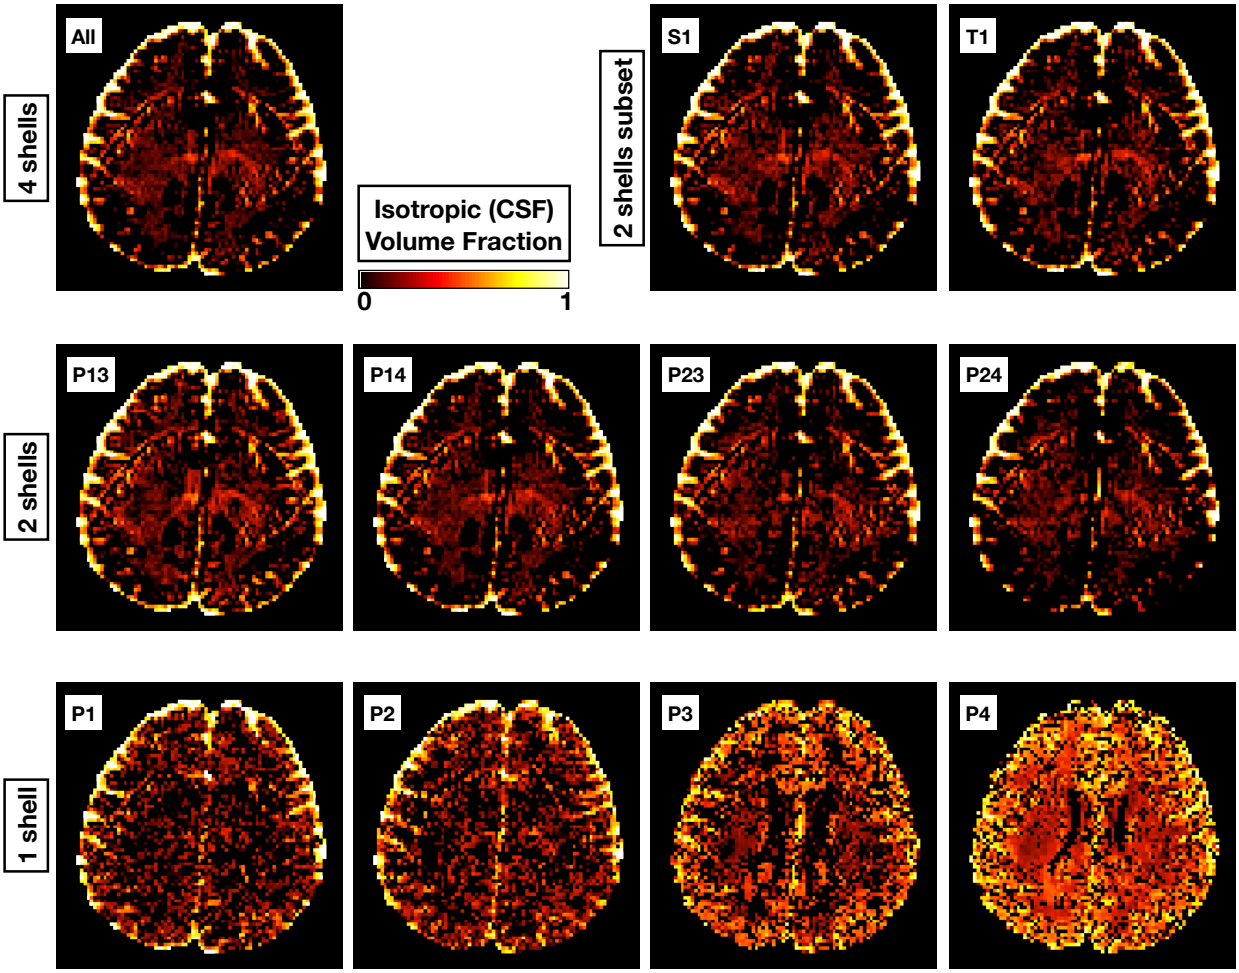

Figure 12: As Fig. 11 but showing maps of isotropic (CSF) volume fraction ν_{iso}.

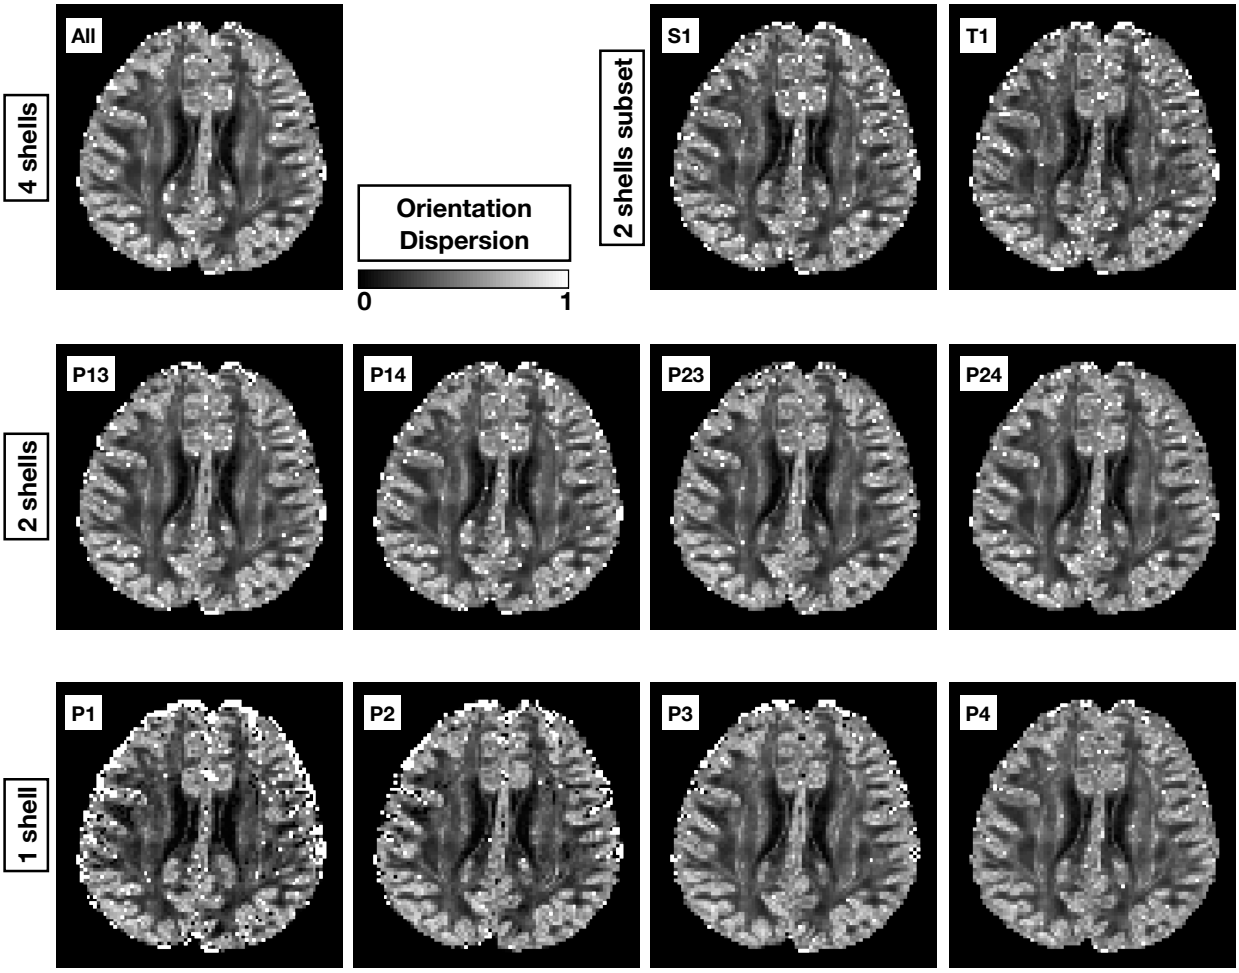

Figure 13: As Fig. 11 but showing maps of orientation dispersion index OD.

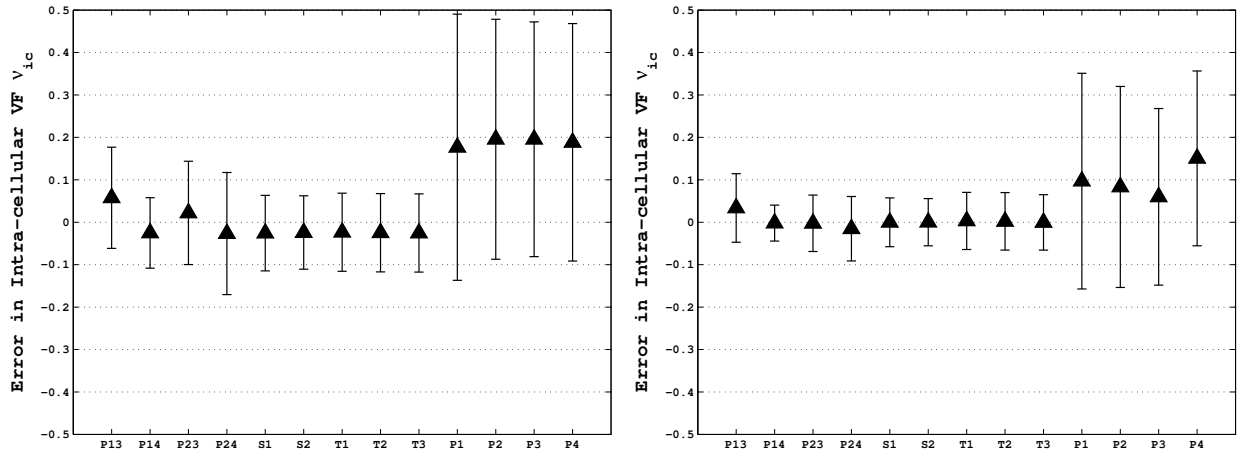

Figure 14: Mean and standard deviation of the estimation errors in ν_{ic} using various protocols when benchmarked against the estimates using the full protocol. The error statistics are calculated separately for gray matter (left) and white matter (right).

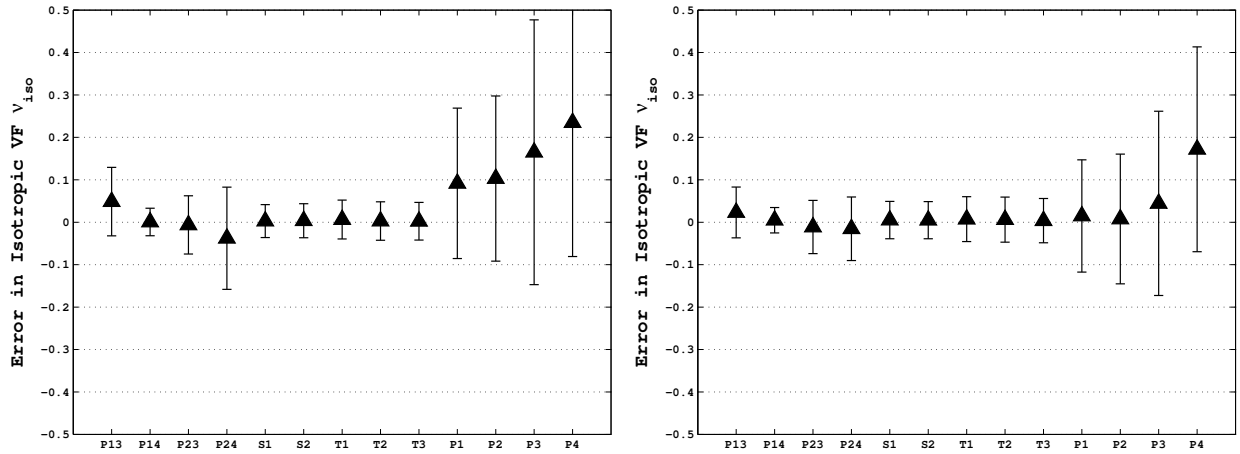

Figure 15: As Fig. 14 but plotting the error statistics of ν_{iso}.

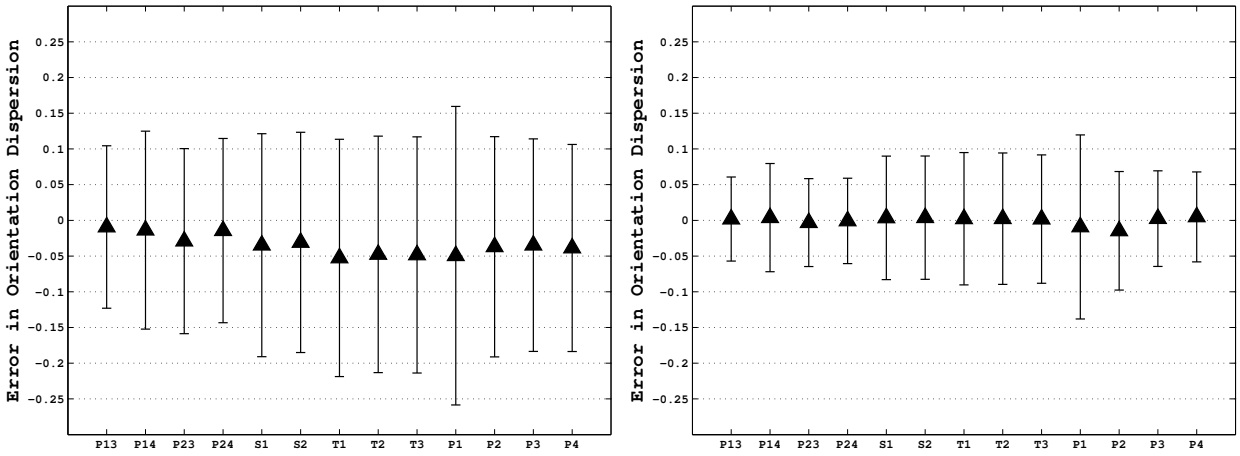

Figure 16: As Fig. 14 but plotting the error statistics of orientation dispersion.

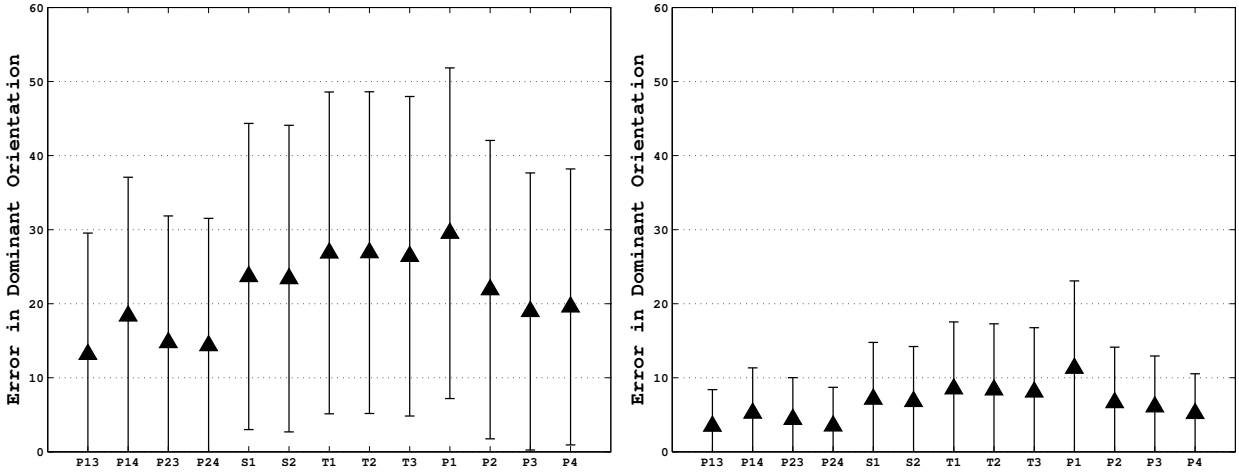

Figure 17: As Fig. 14 but plotting the error statistics of $\boldsymbol{\mu}$.