

Supplementary Information

S.1. Parameter landscape

Figures S.1 and S.2 provide more examples of the topology of the minimization landscape, cf. Fig. 2 from the main text.

S.2. ROI values for the branch selection method

Figure S.3 provides ROI values for the RotInv \pm estimations, as well as for the local branch selection method RotInv ζ . Note that the RotInv ζ values look very similar to those obtained via the prevalence method, cf. Fig. 5 from the main text.

S.3. Noise propagation

Figure S.4 shows results from Monte Carlo simulations of the full MRI protocol (see *Methods*) with 10,000 random combinations of ground truth values uniformly distributed within the biophysically relevant intervals (x -axis, “truth”). The fiber geometry is three identical fiber segments with azimuthal angles $\phi = 0, \pm 2\pi/3$, crossing at an angle $\theta \approx 27^\circ$ with respect to the tract axis, as in Supplementary Figs. S.1 and S.2. Gaussian noise with variance σ^2 is added to both real and imaginary parts of the signal, with absolute value at $b = 0$ normalized to SNR = $1/\sigma$, such that the magnitude signal follows Rician distribution.

Branch degeneracy in Fig. S.4 manifests itself in that branch assignment is not apparent at the level of the rotational invariants — in this case, the moments (panels **a** and **c**) — and becomes evident based on the parameter values (panels **b** and **d**). Practically, $\zeta = +$ corresponds to $D_a > D_e^{\parallel}$ and vice-versa, cf. Eq. (13). In panels **b** and **d**, top row corresponds to parameter estimation based on LEMONADE output, which subsequently served as initialization for the nonlinear fitting of Eq. (7) (middle row), where the LEMONADE branch was pre-selected based on the ground truth values. We can see that noise results in decrease of precision, and that it can accidentally switch the branch. Addition of nonlinear fit (7) notably improves both accuracy and precision relative to LEMONADE.

Note that knowing the branch index ζ beforehand (top two rows in **b** and **d**) makes estimation notably better. Unfortunately, in reality we do not know this index for any given voxel. The 3rd and 4th rows correspond to estimating parameters based on our local branch selection method of Sec. 4.4. Note that the output RotInv ζ is almost as good as that of the (much lengthier) prevalence calculation (bottom row) for sufficiently large SNR (panel **b**), whereas for the lower SNR, spurious parameter values appear (e.g. low f and D_a).

The prevalence method used 100 random initializations within the physically relevant domain of parameters. Generally, intra-axonal parameters f and D_a are more precise than extra-axonal D_e^{\parallel} and D_e^{\perp} ; unfortunately, the branch ratio β is particularly imprecise, prompting the need for “orthogonal” measurements to validate the branch index map in the whole brain.

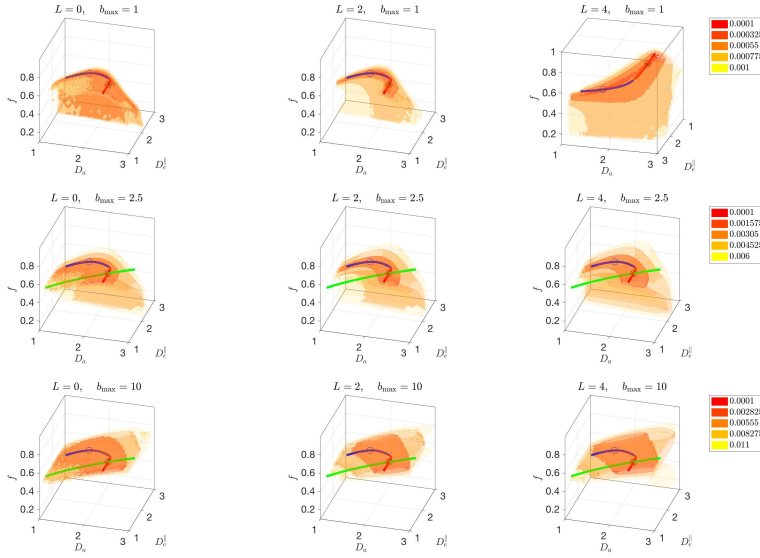


Figure S.1: **Low-energy landscape of the problem (7)**. The F -values are minimized with respect to D_e^\perp and p_l , for the case when the two branches form a single trench within the feasible parameter range. Ground truth values $\{f, D_a, D_e^\parallel, D_e^\perp, p_2\} = \{0.7, 2.4, 1.5, 0.8, 0.7\}$ correspond to three identical fiber segments crossing at an angle $\theta \approx 27^\circ$ to the tract axis. The simulated b -values correspond to those in our human experiments (see Sec. 3.1), with all the 21 b -shells uniformly rescaled to attain the maximal value b_{\max} , such that the bottom row corresponds to the actual acquisition. The two analytical LEMONADE branches (+ red, - blue) match the low-value manifolds, especially for low b_{\max} . Increasing L , the 2-dimensional surface ($L = 0$, corresponding to the two constraints (11a) and (11c) for 4 scalar parameters) gradually turns into 1-dimensional trenches (the full system (11)), while increasing b_{\max} causes flattening of the landscape such that it eventually follows the surface $f/\sqrt{D_a} = \text{const}$ dominated by the intra-axonal water, with the extra-axonal water exponentially suppressed (green line).

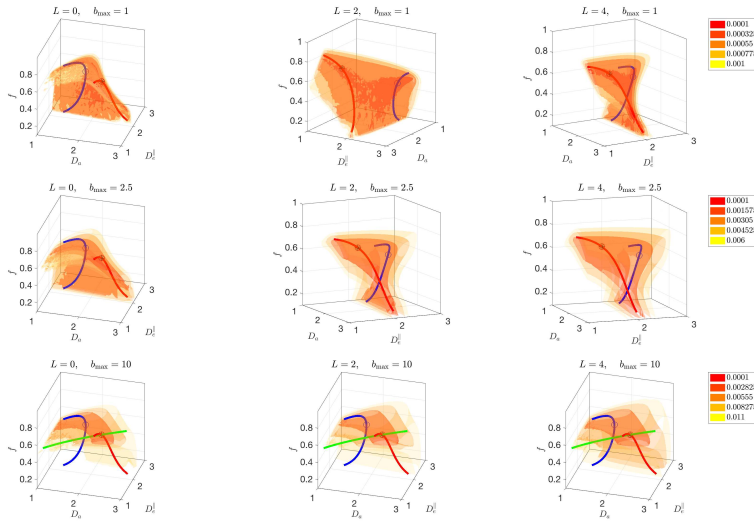


Figure S.2: The same as in Fig. S.1 but with $D_e^\perp = 0.4$. The landscape is highly sensitive to the ground truth values: merely altering one parameter, D_e^\perp , we now have two separate trenches passing through the physically feasible parameter range. They eventually connect (as in Fig. S.1), albeit outside this range. In this case it is particularly easy for spurious minima (e.g. due to noise) to appear in-between the trenches.

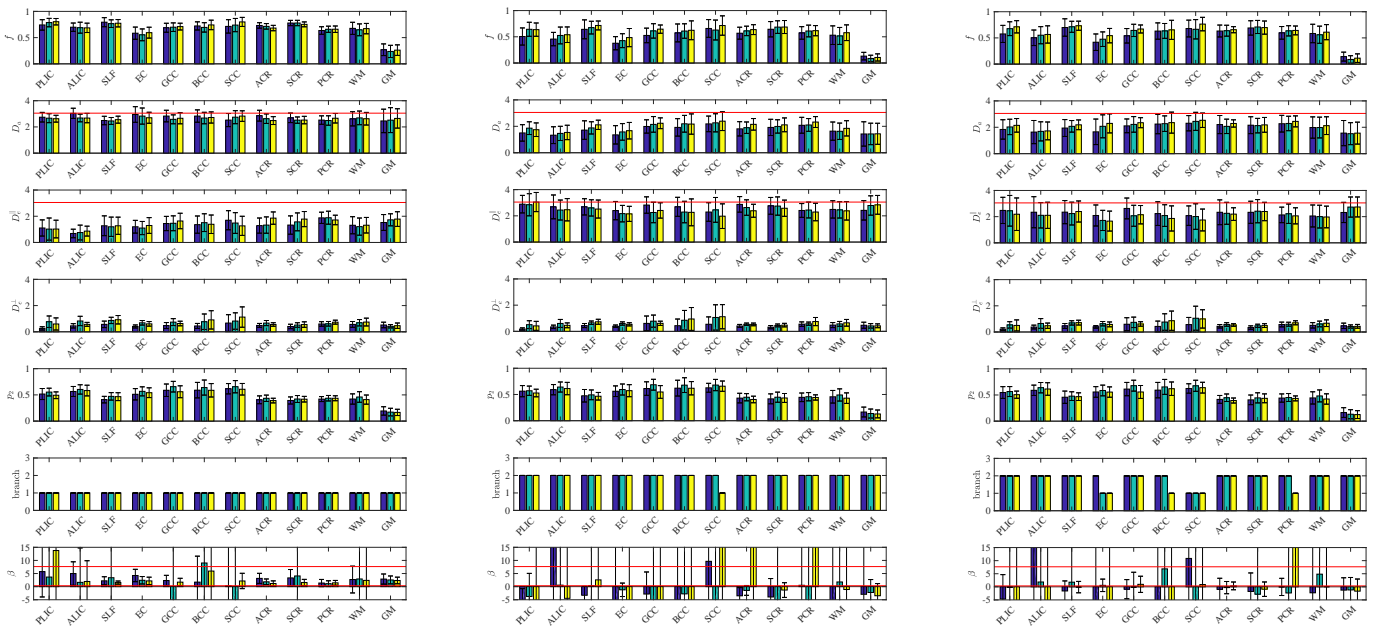

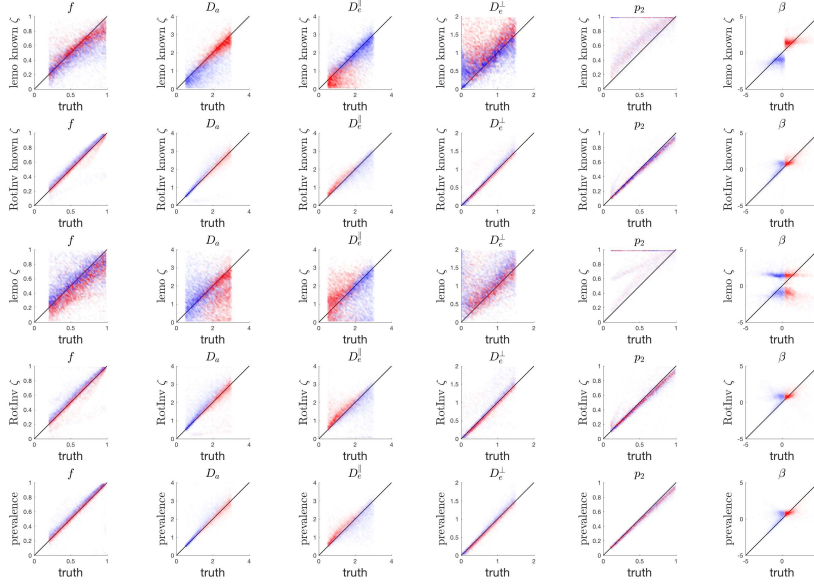



Figure S.3: **ROI values** for the RotInv \pm branches (left and center, correspondingly), as well as for the voxel-wise branch selection method RotInv ζ , cf. Sec. 4.4. Note diffusivity values $D_a > D_e^{\parallel}$ and $D_a < D_e^{\parallel}$ for the \pm branches. Mean values for the RotInv ζ branch selection are quantitatively similar to those based on the prevalence method, Fig. 5, but the error bars are somewhat larger, presumably signifying the fact that our branch selection method is imperfect.



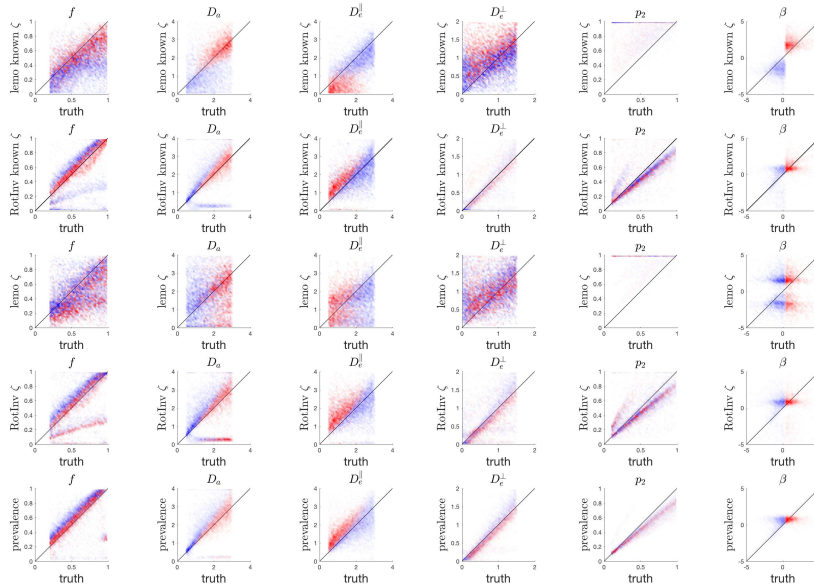
a



b



c



d

Figure S.4: **Noise propagation** for **a, b**: SNR = 100; **c, d**: SNR = 33, at $b = 0$, in estimating moments (panels **a** and **c**) and biophysical parameters (**b** and **d**). Red/blue colors correspond to $\zeta = \pm$ branches assigned based on the ground truth values according to Eq. (13).