

Reducing the Number of Samples in Spatio-Temporal dMRI Acquisition Design

Patryk Filipiak¹ | Rutger Fick¹ | Alexandra Petiet² |
Mathieu Santin² | Anne-Charlotte Philippe² |
Stephane Lehericy² | Philippe Ciuciu³ | Rachid
Deriche¹ | Demian Wassermann^{1,3}

Manuscript word count: 4922

Abstract word count: 184

¹Université Côte d'Azur - Inria Sophia
Antipolis-Méditerranée, France

²CENIR - Center for NeuroImaging
Research, ICM - Brain and Spine Institute,
Paris, France

³Inria, CEA, Université Paris-Saclay, France

Correspondence

Patryk Filipiak, ATHENA Project Team,
Université Côte d'Azur - Inria Sophia
Antipolis-Méditerranée, 2004 Route des
Lucioles, 06902 Valbonne, France
Email: patryk.filipiak@inria.fr

Funding information

the ANR/NSF award NeuroRef; the
European Research Council (ERC) under the
Horizon 2020 research and innovation
program (ERC Advanced Grant agreement
No 694665 : CoBCoM); the MAXIMS grant
funded by ICM's The Big Brain Theory
Program and ANR-10-IAIHU-06; the
programs "Institut des neurosciences
translationnelle" ANR-10-IAIHU-06 and
"Infrastructure d'avenir en Biologie Santé"
ANR-11-INBS-0006.

Purpose: Acquisition time is a major limitation in recovering brain white matter microstructure with diffusion magnetic resonance imaging. The aim of this paper is to bridge the gap between growing demands on spatio-temporal resolution of diffusion signal and the real-world time limitations. The authors introduce an acquisition scheme that reduces the number of samples under adjustable quality loss.

Methods: Finding a sampling scheme that maximizes signal quality and satisfies given time constraints is NP-hard. Therefore, a heuristic method based on genetic algorithm is proposed in order to find sub-optimal solutions in acceptable time. The analyzed diffusion signal representation is defined in the q_T space, so that it captures both spacial and temporal phenomena.

Results: The experiments on synthetic data and in vivo diffusion images of the C57Bl6 wild-type mouse corpus callosum reveal superiority of the proposed approach over random sampling and even distribution in the q_T space.

Conclusion: The use of genetic algorithm allows to find acquisition parameters that guarantee high signal reconstruc-

tion accuracy under given time constraints. In practice, the proposed approach helps to accelerate the acquisition for the use of $q\tau$ -dMRI signal representation.

KEYWORDS

diffusion mri, acquisition design, stochastic optimization

1 | INTRODUCTION

Brain white matter (WM) microstructure recovery with diffusion Magnetic Resonance Imaging (dMRI) requires long acquisition which is unattainable in clinical practice. Dense scanning schemes studied by researchers [3, 8, 57, 58, 59] typically take few hours of imaging time, whereas even one-hour durations are barely acceptable for practitioners. Nonetheless, recent *in vivo* studies of the WM microstructure [2, 7, 16] call for more fine-grained investigation of both space- and time-dependent diffusion. In this work, we aim to bridge the gap between growing demands on spatio-temporal ($q\tau$) resolution of dMRI signal and the real-world time limitations. The above problem can be addressed in at least two ways – either through decreasing the time needed for collecting measurements [10, 20, 36, 53] or through reducing the number of samples [1, 2, 25, 31, 47, 50]. In this study, we use the latter approach.

Time dependence in dMRI has been argued to be an important tool for tissue microstructure analysis [9, 34, 38, 39, 51, 55], accounting either for intra- [3, 4, 5, 45] or extra-cellular diffusion [11, 16, 27, 40, 42]. Nonetheless, a majority of current methods of reconstructing the Ensemble Average Propagator (EAP) [12, 35, 56] do not take into account time decay [17, 26, 44, 57, 58, 59]. Under these circumstances, unified spatio-temporal signal representations ($q\tau$ -dMRI), such as those proposed by Fick et al. [24, 25], are gaining momentum. We believe that this is the moment to grasp such momentum and propose, for the first time, the $q\tau$ -based acquisition design solutions.

The main goal of our study is to find a sampling scheme that maximizes accuracy of a signal representation and satisfies given time constraints. Similar studies were performed for Diffusion Kurtosis Imaging (DKI) [31, 47] or the composite hindered and restricted model of diffusion (CHARMED) [1, 2, 50] including gradient waveform optimization [18, 19]. In this paper, we optimize for the accuracy of $q\tau$ -dMRI signal reconstruction. Additionally, we want our approach to be usable on real-world applications, considering that different acquisition protocols are recommended for different WM microstructure models [52]. For this, we discretize the spatio-temporal search space by performing a dense pre-acquisition of dMRI signal with uniform coverage of the unit sphere, as suggested by Caruyer et al. [13]. Despite discretization, the problem is computationally difficult. It requires selecting an optimal subset of Diffusion Weighted Images (DWIs), which is NP-hard, as we will show in Section 3.

Taking into account that the time complexity of our problem grows exponentially with the domain size, such that global optima cannot be found deterministically within few hours or even few days, we apply a stochastic search engine instead. We use Standard Genetic Algorithm (SGA) [33, 46, 49] for this purpose, which allows us to find approximate solutions in acceptable time.

Our experiments comprise of two granularity levels with respect to the domain size. The coarse-grained level allows us to present the effectiveness of the proposed approach by comparing our results with the global optimum found with lengthy exhaustive search. The fine-grained level outcomes provide us with crucial information about the structure of the optimized acquisition schemes.

We validate our approach on both synthetic diffusion model and real data comprising *in vivo* diffusion images of the C57Bl6 wild-type mouse corpus callosum.

2 | THEORY

In a nutshell, our spatio-temporal signal representation is a cross-product of a 3D space of diffusion gradient directions and a 1D range of diffusion times, hence 4D. This section briefly introduces the mathematical formulation of the suggested representation. For more details we refer to Fick et al. [21, 25, 26].

2.1 | Four-dimensional Ensemble Average Propagator

Let us first define the $q\tau$ -diffusion signal space and its relationship to the 4D Ensemble Average Propagator (EAP) [12, 56]. In dMRI, the EAP $P(\mathbf{R}; \tau)$ describes the probability density that a particle undergoes a displacement $\mathbf{R} \in \mathbb{R}^3$ after diffusion time $\tau > 0$. The EAP is estimated from a set of DWIs, which are obtained by applying two sensitizing diffusion gradients $\mathbf{G} \in \mathbb{R}^3$ of pulse length $\delta > 0$, separated by separation time $\Delta > 0$. Assuming that no diffusion takes place during the pulses ($\delta \rightarrow 0$), the EAP is related to the dMRI signal through the Inverse Fourier Transform (IFT) [56] defined as

$$P(\mathbf{R}; \tau) = \int_{\mathbb{R}^3} E(\mathbf{q}, \tau) e^{i2\pi\mathbf{q}\cdot\mathbf{R}} d\mathbf{q}, \quad (1)$$

where the signal attenuation $E(\mathbf{q}, \tau) = S(\mathbf{q}, \tau)/S_0$, and $S(\mathbf{q}, \tau)$ is the signal measured at the diffusion encoding position $\mathbf{q} \in \mathbb{R}^3$ and diffusion time $\tau = \Delta - \delta/3$, whereas S_0 is the baseline image acquired without diffusion sensitization, i.e. $\mathbf{q} = 0$. We denote $q = |\mathbf{q}|$ and $G = |\mathbf{G}|$, such that $\mathbf{q} = q\mathbf{u}$ and $\mathbf{G} = G\mathbf{v}$ for some 3D unit vectors $\mathbf{u}, \mathbf{v} \in \mathbb{S}^2$. The wave vector \mathbf{q} in Equation (1) is defined as $\mathbf{q} = \gamma\delta\mathbf{G}/2\pi$, where γ is the nuclear gyromagnetic ratio and \mathbf{G} is the applied diffusion gradient vector.

2.2 | $q\tau$ -space signal representation with GraphNet regularization

We reconstruct the continuous EAP from a finite set of DWIs by representing the discretely measured attenuation $E(\mathbf{q}, \tau)$ in terms of the basis coefficients \mathbf{c} of a "Multi-Spherical" 4D $q\tau$ -Fourier basis [21]. The $q\tau$ -basis is formed by the cross-product of a 3D \mathbf{q} -space basis $\Phi_i(\mathbf{q})$ [44] and 1D diffusion time basis $T_j(\tau)$ [25]. The approximated signal attenuation $\hat{E}(\mathbf{q}, \tau, \mathbf{c})$ is given as

$$\hat{E}(\mathbf{q}, \tau, \mathbf{c}) = \sum_{i=1}^{N_q} \sum_{j=1}^{N_\tau} c_{ij} \Phi_i(\mathbf{q}) T_j(\tau) \quad \text{with} \quad \mathbf{c} = [c_{ij}] \in \mathbb{R}^{N_q \times N_\tau}, \quad (2)$$

where N_q and N_τ are the maximum expansion orders of spatial and temporal bases, respectively, and c_{ij} are the weights of the contribution of the i_j^{th} basis function to $\hat{E}(\mathbf{q}, \tau, \mathbf{c})$. As Φ is a Fourier basis over \mathbf{q} , the EAP can be recovered for each voxel as $\hat{P}(\mathbf{R}; \tau, \mathbf{c}) = \text{IFT}_{\mathbf{q}} [\hat{E}(\mathbf{q}, \tau; \mathbf{c})]$.

To estimate \mathbf{c} from a noisy and sparsely sampled $E(\mathbf{q}, \tau)$, we use GraphNet regularization [29], as advised by Fick et al. [25]. Note that the commonly assumed Rician noise perturbing the dMRI signal can be approximated with Gaussian distribution for $\text{SNR} > 5$ [30]. In our GraphNet approach, we use the Laplacian regularization term to compensate for the Gaussian noise and the l_1 -norm regularization to impose sparsity over the basis coefficients. Mathematically, the

above are formulated as follows

$$\begin{aligned} \operatorname{argmin}_{\mathbf{c}} \iint [E(\mathbf{q}, \tau) - \hat{E}(\mathbf{q}, \tau, \mathbf{c})]^2 d\mathbf{q}d\tau + \lambda \iint [\nabla^2 \hat{E}(\mathbf{q}, \tau, \mathbf{c})]^2 d\mathbf{q}d\tau + \alpha \|\mathbf{c}\|_1 \\ \text{subject to } \hat{E}(0, \tau, \mathbf{c}) = 1 \text{ and } \hat{E}(\mathbf{q}, 0, \mathbf{c}) = 1. \end{aligned} \quad (3)$$

The parameters $\lambda, \alpha > 0$ stand for the smoothness and sparsity regularization weights, respectively. We optimize them using five-fold cross-validation.

Having introduced our 4D dMRI signal representation, let us now focus on the acquisition design, which is the main contribution of this paper. Note that our optimization methodology is not strictly related to $q\tau$ -dMRI. In fact, one can replace it with any other approach as long as it allows to recover the original dMRI signal from the subsampled DWIs.

3 | METHODS

We aim to find a sampling scheme that satisfies a given time constraint under adjustable quality loss. For this, we first perform a dense acquisition of dMRI signal. Then, among those densely acquired samples, we seek a fixed-sized subset for which our $q\tau$ -dMRI signal representation reaches the highest accuracy taking the dense acquisition as reference. In the following subsections, we first formulate this goal mathematically as an optimization problem and study its complexity. Later on, we address the problem with a Genetic Algorithm (GA). Then, we define the accuracy measures and describe the experiments.

3.1 | Optimal Acquisition Design

Let $N > 0$ be the number of DWIs in a dense pre-acquisition. As a matter of fact, each DWI corresponds with a certain (q, τ) pair, although for the moment it is more convenient to think of them as an enumerated set of samples. Among those, we want to select a subset of up to $n_{max} < N$ samples, such that the reconstruction accuracy of the dMRI signal is maximized. To this end, we define the objective function $F : \{0, 1\}^N \rightarrow \mathbb{R}$ in the space of binary vectors $x = (x_1, \dots, x_N) \in \{0, 1\}^N$. The assignment $x_i = 1$ for a given $i = 1, \dots, N$ indicates that the i -th sample from the pool of DWIs is included in the subset of interest, whereas $x_i = 0$ determines its exclusion. Formally, our goal is to solve the following optimization problem

$$\begin{aligned} \operatorname{argmin}_x F(x) = \frac{1}{M} \sum_{j=1}^M \|E^{(j)} - \hat{E}_x^{(j)}\|_2^2 \\ \text{subject to } \sum_{i=1}^N x_i \leq n_{max} \quad \text{with } 1 \leq n_{max} \leq N - 1, \end{aligned} \quad (4)$$

where $M > 0$ is the number of voxels in each DWI, $E^{(j)}$ is the normalized signal attenuation captured in the j -th voxel with the $q\tau$ measurements, and $\hat{E}_x^{(j)}$ is the corresponding signal obtained by fitting the $q\tau$ -dMRI representation (defined in Equation 3) to the subsamples indicated by x . Note that the normalization of $E^{(j)}$ is necessary to equalize signal intensities obtained with different (q, τ) parameters.

From now on, we will omit the voxel indexing (j) while referring to E and \hat{E}_x for simplicity. Our goal thus comes down in one sentence: minimize the residuals between E and \hat{E}_x using up to n_{max} samples.

The problem posed in Equation (4), despite its simple formulation, is in fact very difficult to solve analytically. In

Combinatorics, it is known under the name *Knapsack Problem* (KP) and is proved to be NP-hard [32]. Let us remind that the objective of KP is to pick a finite set of items that maximize the total value of the knapsack, while respecting its capacity limitation. In our case, the goal is to select a subset of $q\tau$ -indexed samples that maximize accuracy of dMRI signal recovery, while satisfying given time constraints. Assuming that acquisition time of each DWI is constant, we express the time budget as the maximum number of $q\tau$ -indexed samples in a subset, n_{max} .

The only known deterministic solver of KP is the exhaustive search procedure comprising $N!/(n_{max}!(N-n_{max})!)$ evaluations of the objective function F , which is unacceptable for large N . Thus, we address the problem with a stochastic search engine. For this purpose, we choose SGA [28, 33] due to its reported high performance in large KP instances [14, 37, 41, 43, 54]. Formally, the time complexity of SGA is polynomial, although practically it is difficult to estimate, since certain input parameters, including the number of iterations, need to be chosen experimentally [33].

3.2 | Objective function & performance measures

As defined in Equation (4), the objective of our optimization mechanism is to minimize the residuals between E and \hat{E}_X . We use this quantity as a primary measure of the dMRI signal reconstruction accuracy. Additionally, we are interested in verifying how well the spatio-temporal indices are preserved while using our scheme. To this end, we study a set of commonly used metrics [26, 44], namely

- Return To Origin Probability:

$$\text{RTOP}(\tau) = P(0; \tau) \quad [1/\text{m}^3]$$

- Return To Axis Probability:

$$\text{RTAP}(\tau) = \int_{\mathbb{R}} P(R\mathbf{r}_{\parallel}; \tau) dR \quad [1/\text{m}^2]$$

- Return To Plane Probability:

$$\text{RTPP}(\tau) = \int_{\mathbb{R}} \int_{\{\mathbf{r} \in \mathbb{S}^2 : \mathbf{r} \cdot \mathbf{r}_{\parallel} = 0\}} P(R\mathbf{r}_{\perp}; \tau) d\mathbf{r}_{\perp} dR \quad [1/\text{m}]$$

- Mean Squared Displacement:

$$\text{MSD}(\tau) = \int_{\mathbb{R}} \int_{\mathbb{S}^2} P(R\mathbf{r}; \tau) R^2 d\mathbf{r} dR \quad [\text{m}^2]$$

for a given displacement $\mathbf{R} = R\mathbf{r}$, as defined in Section 2. Let us mention that two of the above metrics, i.e. RTAP and RTPP, assume that WM is modeled by parallel cylinders with the vectors \mathbf{r}_{\parallel} parallel and \mathbf{r}_{\perp} perpendicular to the cylinder axis.

Granularity levels:

The acquisition schemes that we study are 800-dimensional binary feature vectors. Dealing with such high-dimensional data involves the risk of bias towards a dense pre-acquisition scheme or over-fitting to dMRI signal noise. In order to avoid these pitfalls, we consider the following two granularity levels:

- *Optimization by shells* – where each of the shells in our multishell dense acquisition is either taken as a whole or discarded completely. In this scenario, the task is to subsample among 5 available τ values and 8 available G values, which mimics the experiment described by Alexander et al. [1, 2]. Particularly, for the budget size $n_{max} = 100$, the search space is reduced to 658,008 possible solutions, hence we are able to find the global optimum of the problem by performing exhaustive search and use it as a reference (cf. Figure 1). Our proposed optimizer uses a set of 100 candidate solutions per iteration in this scenario, which is a commonly used setting for SGA [6].
- *Optimization by measures* – where solutions are picked from all the 800 variants of densely acquired samples. Considering relatively high complexity of this scenario, we chose experimentally the population size of our SGA optimizer to be 1000 candidate solutions per iteration.

3.3 | Experiments setup

In the experiments, we analyze both synthetic and real diffusion data using the protocol described below.

Dense pre-acquisition:

Our initial dense pre-acquisition covered 40 shells, each of which comprised 20 directions and one b_0 -image, i.e. $40 \times 20 = 800$ DWIs plus 40 non-weighted images. We used combinations of 8 gradient strengths $G \in \{50, 100, 150, 200, 250, 300, 350, 400\}$ [mT/m] (magnitudes available in pre-clinical scanners) and 5 separation times $\Delta \in \{10.8, 13.1, 15.4, 17.7, 20.0\}$ [ms] to parametrize our (q, τ) acquisition space. The corresponding b-values ranged from 41 to 5248 s/mm^2 . The gradient duration $\delta = 5$ ms remained constant throughout the experiments. In each of the $8 \times 5 = 40$ shells, we followed the uniform distribution of directions suggested by Caruyer et al. [13].

Time constraints:

We considered four variants of time limits expressed as budget sizes $n_{max} \in \{100, 200, 300, 400\}$ out of 800 DWIs. The remaining 40 b_0 -images were excluded from the optimization domain, as they were used by default in every acquisition scheme.

Other approaches:

We compared our method with two alternative sampling schemes. One of them, called *random*, used the uniform random distribution of $q\tau$ samples in the index space $\{1, \dots, N\}$. In the second one, referred to as *even*, we picked each i -th sample for $i = \lfloor kN/n_{max} \rfloor$ and $k = 1, \dots, n_{max}$. Considering that the samples were ordered by G and Δ , the *even* subsampling scheme ensured nearly equal number of samples per shell, i.e. $\lfloor n_{max}/40 \rfloor$ or $\lceil n_{max}/40 \rceil$.

3.4 | Diffusion data

We used the following two data sets:

Synthetic data:

As first data set, we generated time-dependent diffusion data using our Python-based [22, 23] implementation of the two-compartment model with the intra-cellular fraction modeled with a set of Watson dispersed cylinders [12, 60] and the extra-cellular one modeled with a temporal zeppelin [11]. The Watson distribution is an antipodally symmetric distribution, centered around the principal orientation μ , describing a density controlled by concentration parameter $\kappa > 0$, which is inversely related to the dispersion of cylinders [12, 60]. Ronen et al. [48] modeled the axon dispersion in

the corpus callosum with a spherical Gaussian distribution and found out that typical values of its standard deviation vary between 10 and 25 degrees, which translates to the interval $3 < \kappa < 16$. In another study, Zhang et al. [60] used the values of $\kappa \in \{0, 0.25, 1, 4, 16\}$. In our experiment, we chose $\kappa = 4$ from the intersection of both sets.

Table 1 summarizes all the model's parameters that we took from Burcaw et al. [11], Alexander et al. [2], and Zhang et al. [60].

[Table 1 about here.]

Apart from the original noiseless signal, we also studied the two variants of the signal with incorporated Rician noise, having respective Signal-to-Noise Ratio (SNR) set to 20 (13dB) and 10 (10dB).

Real data:

In the second data set, we used in vivo diffusion images of the corpus callosum of C57Bl6 wild-type mouse. All animal experiments were performed in accordance with the EU Directive 2010/63/EU for animal experiments. Obtaining the initial dense pre-acquisition took approximately 2h10min on an 11.7 Tesla Bruker scanner (Bruker Biospec 117/16 USR horizontal bore, 750 mT/m gradients, Paravision 6.0.1, Ettlingen, Germany). The data consists of $96 \times 160 \times 12$ voxels covering a field of view $110 \times 110 \times 500 \mu\text{m}^3$ each. The average SNR of the images is $10\text{dB} \pm 2$. We manually created a brain mask and corrected the data from eddy currents and motion artifacts using FSL's *eddy*.

4 | RESULTS

Exhaustive search:

As we mentioned earlier, we defined the coarse-grained *optimization by shells* scenario in such a way that it could be solved with exhaustive search and serve as a reference. The running time of a parallel code for this approach was about one week. Let us mention that using the same exhaustive search implementation for the *optimization by measures* scenario would take more than 10^{200} centuries.

Figure 1 illustrates all the 658,008 feasible acquisition schemes in the *optimization by shells* scenario for the in silico experiment with $n_{max} = 100$. The schemes are arranged from best to worst. The top-left plot shows the mean squared errors (MSEs) of signal reconstruction, whereas the top-right one depicts the normalized Hamming distances from the global optimum ± 1 standard deviation. In order to visualize the analyzed (G, Δ) parameter space, the percentiles $pc = 0\%, 1\%, 10\%, 50\%, 90\%$ are annotated on both plots, showing respectively the global optimum, the top 1% solutions, the top 10% solutions, etc. The cumulative averages of acquisition schemes, corresponding with these percentiles, are shown in the heat maps at the bottom. The colors reflect the likelihood of a given (G, Δ) pair in the scheme. The heat maps (a) and (b) represent, respectively, the global optimum and its proximity. The interval of percentiles between $pc = 10\%$ and $pc = 90\%$, as shown in the heat maps (c)-(e), contains a spectrum of feasible acquisition schemes with similar MSEs and almost equally large Hamming distances from the global optimum.

Note the arrangements of high likelihood (red-colored squares) in Figure 1(a). There is one red field in the area of highest G-values located in the middle of the Δ range. Then, there are four other red fields spread evenly across G- and Δ parameter spaces. The pattern gradually disappears as we move away from the global optimum. Starting from about 10th percentile, the differences in intensities on the heat map slowly fade.

[Figure 1 about here.]

In silico experiments:

As shown in Figure 2, convergence of SGA in the *optimization by shells* scenario is apparently reached before the 30th iteration. Based on that observation, we cautiously chose the termination condition for all experiments to be the fixed threshold of 40 iterations, thus leaving a safety margin of 10 iterations.

[Figure 2 about here.]

Figure 3 (left-column plots) and Table 2 summarize the results of the tested subsampling schemes — *ours*, *random*, and *even* — for the in silico experiment, averaged over 30 runs. In the rows of Table 2, we present respectively the noiseless case and the two variants with incorporated Rician noise having SNR = 20 (13dB) and 10 (10dB). Within each case, we consider four time budget limits $n_{max} \in \{100, 200, 300, 400\}$. The scores are expressed with normalized root mean squared errors (NRMSEs) and the corresponding standard deviations (STDs). The STD is omitted in the *even* scheme which is deterministic and thus its STD=0.

[Figure 3 about here.]

As Table 2 shows, our approach outperforms the other two in all studied cases, in both granularity levels, reaching lowest error values and standard deviations. It is also worth noticing that the addition of Rician noise induces overfitting of the q_T -dMRI signal representation for $n_{max} > 200$. As a consequence, our approach handles better the noisy data in the *optimization by shells* scenario, having 20 times less free parameters, than the *optimization by measures*.

We compared all the pairs of results, i.e. *ours* vs. *random* and *ours* vs. *even*, in both granularity levels, using paired two-sample Student's t -tests with the Bonferroni adjusted significance level $\alpha = 10^{-5}$ and the number of degrees of freedom $2n - 2 = 58$. In each case, our approach was statistically significantly better than the other two schemes.

[Table 2 about here.]

In order to study the stability of our proposed method, we distilled a single best acquisition scheme found in each of the 30 runs of SGA. Next, we arranged these solutions by MSEs and selected the top 10% out of them. Figure 4 illustrates the averages of those schemes. Note the perfect similarity between the best schemes found by SGA, presented in Figure 4(a), and the global optimum depicted in Figure 1(a). Furthermore, the concentrations of high likelihood (red squares) in Figure 4 tend to form shapes that are especially visible in the *optimization by shells* scenario. The latter suggests that SGA repeatedly converges to the same or highly similar solutions in each run, thus giving stable outputs.

[Figure 4 about here.]

In vivo experiments:

Figure 3 (right-column plots) and Table 3 present the results for the C57Bl6 wild-type mouse corpus callosum (CC). Table 3 details the three regions of interest, i.e. *genu*, *body*, and *splenium*, and the time budgets $n_{max} \in \{100, 200, 300, 400\}$. Our method outperforms the other two significantly, which we verified using the same type of Student's t -test as above. However, unlike in silico experiment, there is no superiority of *optimization by shells* over *optimization by measures* here. In fact, the residuals in both scenarios are comparable, with slightly lower averages in *optimization by measures* for small n_{max} , whereas *optimization by shells* gives lower averages for large n_{max} .

[Table 3 about here.]

Analogously to the in silico experiment, we illustrate the averaged top 10% acquisition schemes obtained with SGA in the form of heat maps depicted in Figure 5. Here also the gradually smoothing patterns are present as in Figure 4.

[Figure 5 about here.]

Spatio-temporal indices:

Figures 6 and 7 present the reconstruction of spatio-temporal indices RTOP, RTAP, RTPP, $MSD \pm 1$ standard deviation with $n_{max} = 100$, obtained for in silico and in vivo experiments respectively. The black lines show the reference curves obtained from the densely acquired signal, whereas the color lines represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes. In addition, Supporting Information Figures S1 and S2 present the diffusion tensor related curves of mean, axial and radial diffusivity, and fractional anisotropy for in silico and in vivo experiments respectively.

As it is seen in Figure 6, all the indices are well preserved under subsampling by the three tested schemes in the noiseless case of in silico experiment. However, the discrepancies between the reconstructed and the reference curves readily increase with the addition of noise. Also, index dispersion grows as SNR decreases and diffusion time increases.

The differences between the reconstructed and the reference curves are similar in the in vivo experiment presented in Figure 7. Generally, all the tested schemes give comparable outputs, although the results of our approach are less dispersed in this case. The highest accuracy is reached at the *splenium* region, while there are considerable errors in the *body*.

[Figure 6 about here.]

[Figure 7 about here.]

5 | DISCUSSION

It is beyond doubt that acquisition time matters, especially in long-lasting protocols like dMRI. In this paper, we propose a method to shorten acquisition time with reduced signal loss by optimizing the sampling scheme in the (q, τ) parameter space. Additionally, we want our scheme to preserve temporal changes in diffusion signal, hence we study the $q\tau$ -indices presented in Section 3.

In practice, our approach helps to accelerate the acquisition for the use of $q\tau$ -dMRI signal representation. It is crucial, since the latter allows to extrapolate the signal outside $q\tau$ -samples.

5.1 | Our approach maximizes signal reconstruction accuracy

The accuracy of WM microstructure recovery grows monotonically with an increase of dMRI acquisition density, although the contributions of particular DWIs to the measured signal attenuation are not equal. This indicates the need for identifying a variable-density for acquiring the dMRI signal. Our method, described in Section 3, seeks a fixed-sized set of samples that provide the most accurate signal reconstruction with respect to the reference data obtained with the dense acquisition scheme. In other words, our approach allows to reach high accuracy within a controllable time budget.

As we can see, by comparing Figures 1 and 4, our approach is able to find nearly-optimal solutions in less than 30 iterations of the proposed algorithm (cf. Figure 2). The accuracy of dMRI signal reconstruction obtained with our technique is significantly better than *random* or *even* subsampling of the densely acquired signal, as we can observe in Tables 2 and 3, as well as Figure 3

5.2 | Scheme optimization is most profitable in short acquisition times

In our experiments, we studied the four time budgets $n_{max} \in \{100, 200, 300, 400\}$ out of 800 densely acquired samples. Both, in silico (cf. Table 2 and Figure 3) and in vivo experiments (cf. Table 3 and Figure 3), showed that superiority of

our technique was most evident in the cases with tightest time constraints $n_{max} = 100$, which is $1/8$ of the original time span. For $n_{max} \geq 200$, the differences in residuals of the dMRI signal reconstruction among the three tested approaches decreased gradually. Also, overfitting of the $q\tau$ -dMRI signal emerged as n_{max} grew. This means that the less imaging time we had, the more we benefited from the acquisition scheme optimization. On the other hand, when small speed-ups were concerned (e.g. $1/2$ of the original time span), our optimization seemed unnecessary. The way to go in such cases was as simple as to remove gradient directions from the dMRI scanning protocol evenly, i.e. every second one for $1/2$ time save, every third one for $1/3$, and so on.

5.3 | The optimized schemes form repetitive patterns

In agreement with Caruyer et al. [13], we believe there can be no single all-purpose optimal acquisition scheme. The imaging parameters depend on numerous factors including the needs of a given WM microstructure recovery method, the time constraints, and the physical limitations of a scanner. That is why we propose a data-driven mechanism which is able to adapt to varying imaging conditions by maximizing the accuracy of the signal reconstructed with $q\tau$ -dMRI.

Alexander [1] and Caruyer et al. [13] postulated that an angular coverage of gradient directions across the sphere should be “as uniform as possible” in order to ensure orientational invariance of a scheme. The results of our study partly comply with those remarks. Note that the additional parameters which controlled the distribution of gradient directions introduced in the *optimization by measures* scenario improved the reconstruction accuracy in the noiseless cases only, compared to the coarse-grained *optimization by shells*. Apparently, the ad hoc alterations that we made possible after freeing the space of gradient directions were of little importance for the signal reconstruction accuracy. We thus conclude that studying our problem at such a fine-grained level might be an overhead in real-world applications. On the other hand, there is a possibility that choosing a different (i.e. non-uniform) scheme for the dense pre-acquisition would lead to slightly different conclusions in this regard, which leaves a room for further improvement of the proposed technique.

Apparently, setting up the parameters G and Δ is more difficult than handling gradient directions. An analogue of the simplistic “as uniform as possible” strategy, that we had called the *even* scheme in our experiments, turned out insufficient in this case. Tables 2 and 3 show that the above approach gave mediocre results, roughly comparable to the naive *random* subsampling scheme.

The heat maps presented in Figures 4 and 5 shed some light on the optimized schemes obtained with our approach. The distribution of parameters forms clusters of points in the (G, Δ) space, especially in the *optimization by shells* scenario. In the *optimization by measures*, both in silico and in vivo, the dominance of low G -values emerges as n_{max} increases.

In the similar experiment described by Alexander [1], the distribution of pulse lengths Δ was close to uniform, although it came with the dominance of largest G -values which is not the case in our study. The above discrepancy is most probably caused by the differences in objective functions. Alexander’s work targeted particular biomarkers like axon density and radius in WM, whereas we maximize reconstruction accuracy of a dMRI signal using our $q\tau$ signal representation. This in turn brings us again to the earlier remark that the notion of acquisition scheme optimality is highly problem-dependent.

5.4 | Globally optimal acquisition schemes may be beyond reach

It is proved that the SGA optimization engine, that we use here, will always converge to global optimum when given *enough time* [15]. Nevertheless, the question “how much time is enough?” is generally unanswerable, as it requires

thorough analysis of each problem in hand.

For better understanding of our parameter space, we illustrate the full spectrum of feasible acquisition schemes at the top row of Figure 1. The solutions are arranged from best to worst, i.e. from the lowest MSE of signal reconstruction on the left-hand side boundary to the highest one on the right. Such arrangements is convenient to visualize the accuracy of potential solutions, although we must point out that it disregards the space topology. On the top-left plot in Figure 1, there is a narrow valley in the proximity of global optimum and the high peek on the other side of the plot. Between them, there is a large interval of low steepness. Such a curve shape suggests that the globally best solution can be difficult to find using randomized approaches, since it is not surrounded by a considerable neighborhood of nearly-optimal ones. Indeed, the curve of the normalized Hamming distance to the global optimum (in the feature space), depicted at the top-right plot in Figure 1, grows rapidly with the quantiles pc . To make matters worse, the low steepness in the middle of the top-left plot indicates a possibly large number of local minima or plateau areas which in turn would hinder the numerical optimization techniques like GAs. Both those observations suggest that with an increase of problem complexity, chances of finding global optima will decrease rapidly.

Our experiments showed that SGA was able to find the global optimum in the *optimization by shells* scenario with synthetic diffusion data and $n_{max} = 100$ (cf. Figures 1 and 4). However, it is very likely that in difficult problems like *optimization by measures*, we found sub-optimal solutions instead.

5.5 | Our approach preserves spatio-temporal indices when reducing acquisition time

As it is seen in Figures 6 and 7, the four analyzed spatio-temporal indices are preserved when using our approach. The *even* subsampling scheme produces comparably good averages, however the curves are much more dispersed. On the other hand, the *random* subsampling scheme often gives closer approximations, albeit disregarding the shape of the curve producing a linear fit. Our approach better reproduces the shape of the reference curves although with an increased bias. Additionally, our method performs better on the diffusion tensor indices, particularly Fractional Anisotropy, presented in Supporting Information Figures S1 and S2.

The observed agreement with the reference $q\tau$ indices, which were not included in the objective function introduced in Section 3, suggests that our method is able to generalize well. With $n_{max} = 100$, we avoided overfitting to noise that was present in the densely pre-acquired signal. On the other hand, it might be interesting to extend the objective function, by adding multiple criteria covering, for instance, mean squared displacement or fractional anisotropy.

We also find it promising that the spatio-temporal signal representation itself turned out to harmonize efficiently with our approach. Note that all the plots of $q\tau$ indices presented here illustrate the results for the tightest time budget $n_{max} = 100$. This implies again that a great deal of acquisition time can be saved without much compromise on dMRI signal accuracy.

The studied diffusion times, 10-20ms, cover the range of Δ values that were feasible for the pulsed-gradient spin echo mode due to the T2 decay time. Note that the $q\tau$ indices computed with our approach smoothly interpolate between sampled diffusion times. Also, our $q\tau$ -dMRI signal representation is able to extrapolate the index values outside the measured time interval, although the accuracy in such cases requires further investigation.

6 | CONCLUSIONS

We proposed the spatio-temporal dMRI acquisition design that greatly reduces the number of $q\tau$ samples under the adjustable quality loss. Despite the fact that selecting a sampling scheme that maximizes brain white matter

reconstruction accuracy and satisfies given time constraints is NP-hard, our stochastic optimization mechanism based on genetic algorithm found sub-optimal solutions efficiently.

The experiments on both synthetic diffusion data and in vivo images of the C57Bl6 wild-type mouse corpus callosum revealed superiority of our technique over random subsampling and even distribution in the $q\tau$ space. Our approach performed best under the tightest among all the considered time constraints, leading to reduction of acquisition time to $1/8$ of the original time span. Additionally, we observed repetitive patterns in our optimized schemes that gave us crucial insight about the space of acquisition parameters. The latter is really promising for future identification of optimal sampling density.

In this study, we assumed availability of a densely acquired dMRI signal for reference, although it is not often the case. Nonetheless, we believe that our preliminary work will allow to uncover the optimal sampling density in the (q, τ) space for optimal recovery of the EAP, and thus the optimal range of acquisition parameters G and Δ .

Future work should target the reproducibility of our approach among different subjects and scanners. For instance, the risk of peripheral nerve stimulation while using high gradient magnitudes should be taken into account. Also, different areas of WM tissue need to be studied, including for instance crossing fibers. Finally, the optimizer itself might be improved to ensure faster convergence and adaptability, and thus achieve lower average quality loss of solutions.

ACKNOWLEDGEMENTS

This work has received funding from the ANR/NSF award NeuroRef; the European Research Council (ERC) under the Horizon 2020 research and innovation program (ERC Advanced Grant agreement No 694665 : CoBCoM); the MAXIMS grant funded by ICM's The Big Brain Theory Program and ANR-10-IAIHU-06; the programs "Institut des neurosciences translationnelle" ANR-10-IAIHU-06 and "Infrastructure d'avenir en Biologie Santé" ANR-11-INBS-0006.

REFERENCES

- [1] Alexander DC. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magnetic Resonance in Medicine* 2008;60(2):439–448.
- [2] Alexander DC, Hubbard PL, Hall MG, Moore EA, Ptito M, Parker GJ, et al. Orientationally invariant indices of axon diameter and density from diffusion MRI. *Neuroimage* 2010;52(4):1374–1389.
- [3] Assaf Y, Blumenfeld-Katzir T, Yovel Y, Basser PJ. AxCaliber: a method for measuring axon diameter distribution from diffusion MRI. *Magnetic resonance in medicine* 2008;59(6):1347–1354.
- [4] Assaf Y, Cohen Y, et al. In vivo and in vitro bi-exponential diffusion of N-acetyl aspartate (NAA) in rat brain: A potential structure probe? *NMR in biomedicine* 1998;11(2):67–74.
- [5] Assaf Y, Freidlin RZ, Rohde GK, Basser PJ. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic Resonance in Medicine* 2004;52(5):965–978.
- [6] Baluja S, Caruana R. Removing the genetics from the standard genetic algorithm. In: *Machine Learning: Proceedings of the Twelfth International Conference*; 1995. p. 38–46.
- [7] Bar-Shir A, Avram L, Özarslan E, Basser PJ, Cohen Y. The effect of the diffusion time and pulse gradient duration ratio on the diffraction pattern and the structural information estimated from q-space diffusion MR: experiments and simulations. *Journal of Magnetic Resonance* 2008;194(2):230–236.

- [8] Barazany D, Basser PJ, Assaf Y. In vivo measurement of axon diameter distribution in the corpus callosum of rat brain. *Brain* 2009;132(5):1210–1220.
- [9] Beaulieu C, Allen PS. An in vitro evaluation of the effects of local magnetic-susceptibility-induced gradients on anisotropic water diffusion in nerve. *Magnetic resonance in medicine* 1996;36(1):39–44.
- [10] Boyer C, Chauffert N, Ciuciu P, Kahn J, Weiss P. On the generation of sampling schemes for magnetic resonance imaging. *SIAM Journal on Imaging Sciences* 2016;9(4):2039–2072.
- [11] Burcaw LM, Fieremans E, Novikov DS. Mesoscopic structure of neuronal tracts from time-dependent diffusion. *NeuroImage* 2015;114:18–37.
- [12] Callaghan PT. Pulsed-gradient spin-echo NMR for planar, cylindrical, and spherical pores under conditions of wall relaxation. *Journal of magnetic resonance, Series A* 1995;113(1):53–59.
- [13] Caruyer E, Lenglet C, Sapiro G, Deriche R. Design of multishell sampling schemes with uniform coverage in diffusion MRI. *Magnetic resonance in medicine* 2013;69(6):1534–1540.
- [14] Chu PC, Beasley JE. A genetic algorithm for the multidimensional knapsack problem. *Journal of heuristics* 1998;4(1):63–86.
- [15] Davis TE, Principe JC. A Markov chain framework for the simple genetic algorithm. *Evolutionary computation* 1993;1(3):269–288.
- [16] De Santis S, Jones DK, Roebroek A. Including diffusion time dependence in the extra-axonal space improves in vivo estimates of axonal diameter and density in human white matter. *NeuroImage* 2016;130:91–103.
- [17] Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. Regularized, fast, and robust analytical Q-ball imaging. *MR in Medicine* 2007;58(3):497–510.
- [18] Drobnyak I, Alexander DC. Optimising time-varying gradient orientation for microstructure sensitivity in diffusion-weighted MR. *Journal of Magnetic Resonance* 2011;212(2):344–354.
- [19] Drobnyak I, Siow B, Alexander DC. Optimizing gradient waveforms for microstructure sensitivity in diffusion-weighted MR. *Journal of Magnetic Resonance* 2010;206(1):41–51.
- [20] Feinberg DA, Setsompop K. Ultra-fast MRI of the human brain with simultaneous multi-slice imaging. *Journal of magnetic resonance* 2013;229:90–100.
- [21] Fick R, Petiet A, Santin M, Philippe AC, Lehericy S, Deriche R, et al. Multi-Spherical Diffusion MRI: Exploring Diffusion Time Using Signal Sparsity. In: *MICCAI 2016 Workshop on Computational dMRI (CDMRI'16)*; 2016..
- [22] Fick R, Wassermann D, Deriche R, Dmipy: An Open-source Framework for Reproducible dMRI-Based Microstructure Research (Version 0.1). Zenodo.; 2018. <http://doi.org/10.5281/zenodo.1188268>.
- [23] Fick R, Wassermann D, Deriche R. Mipy: An Open-Source Framework to improve reproducibility in Brain Microstructure Imaging. In: *OHBM 2018-Human Brain Mapping*; 2018. p. 1–4.
- [24] Fick R, Wassermann D, Pizzolato M, Deriche R. A unifying framework for spatial and temporal diffusion in diffusion MRI. In: *International Conference on Information Processing in Medical Imaging Springer*; 2015. p. 167–178.
- [25] Fick RH, Petiet A, Santin M, Philippe AC, Lehericy S, Deriche R, et al. Non-parametric graphnet-regularized representation of dMRI in space and time. *Medical Image Analysis* 2018;43:37–53.
- [26] Fick RH, Wassermann D, Caruyer E, Deriche R. MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data. *NeuroImage* 2016;134:365–385.

- [27] Fieremans E, Burcaw LM, Lee HH, Lemberskiy G, Veraart J, Novikov DS. In vivo observation and biophysical interpretation of time-dependent diffusion in human white matter. *NeuroImage* 2016;129:414–427.
- [28] Goldberg DE. Genetic algorithms in search, optimization, and machine learning, 1989. Reading: Addison-Wesley 1989;.
- [29] Grosnick L, Klingenberg B, Katovich K, Knutson B, Taylor JE. Interpretable whole-brain prediction analysis with GraphNet. *NeuroImage* 2013;.
- [30] Gudbjartsson H, Patz S. The Rician distribution of noisy MRI data. *Magnetic resonance in medicine* 1995;34(6):910–914.
- [31] Hansen B, Lund TE, Sangill R, Jespersen SN. Experimentally and computationally fast method for estimation of a mean kurtosis. *Magnetic resonance in medicine* 2013;69(6):1754–1760.
- [32] Hochbaum DS. Approximation algorithms for NP-hard problems. PWS; 1996.
- [33] Holland JH. Adaptation in natural and artificial systems. An introductory analysis with application to biology, control, and artificial intelligence. Ann Arbor, MI: University of Michigan Press 1975;.
- [34] Horsfield MA, Barker GJ, McDonald WI. Self-diffusion in CNS tissue by volume-selective proton NMR. *Magnetic resonance in medicine* 1994;31(6):637–644.
- [35] Kärgler J, Heink W. The propagator representation of molecular transport in microporous crystallites. *Journal of Magnetic Resonance (1969)* 1983;51(1):1–7.
- [36] Keil B, Blau JN, Biber S, Hoecht P, Tountcheva V, Setsompop K, et al. A 64-channel 3T array coil for accelerated brain MRI. *Magnetic resonance in medicine* 2013;70(1):248–258.
- [37] Khuri S, Bäck T, Heitkötter J. The zero/one multiple knapsack problem and genetic algorithms. In: Proceedings of the 1994 ACM symposium on Applied computing ACM; 1994. p. 188–193.
- [38] Kunz N, Sizonenko SV, Hüppi PS, Gruetter R, Looij Y. Investigation of field and diffusion time dependence of the diffusion-weighted signal at ultrahigh magnetic fields. *NMR in biomedicine* 2013;26(10):1251–1257.
- [39] Latour LL, Svoboda K, Mitra PP, Sotak CH. Time-dependent diffusion of water in a biological model system. *Proceedings of the National Academy of Sciences* 1994;91(4):1229–1233.
- [40] Lee HH, Fieremans E, Novikov DS. What dominates the time dependence of diffusion transverse to axons: Intra-or extra-axonal water? *NeuroImage* 2017;.
- [41] Michalewicz Z, Arabas J. Genetic algorithms for the 0/1 knapsack problem. In: International Symposium on Methodologies for Intelligent Systems Springer; 1994. p. 134–143.
- [42] Novikov DS, Jensen JH, Helpert JA, Fieremans E. Revealing mesoscopic structural universality with diffusion. *Proceedings of the National Academy of Sciences* 2014;111(14):5088–5093.
- [43] Olsen AL. Penalty functions and the knapsack problem. In: Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on IEEE; 1994. p. 554–558.
- [44] Özarlan E, Koay CG, Shepherd TM, Komlosh ME, İrfanoğlu MO, Pierpaoli C, et al. Mean apparent propagator (MAP) MRI: A novel diffusion imaging method for mapping tissue microstructure. *NeuroImage* 2013;78:16–32.
- [45] Palombo M, Ligneul C, Najac C, Le Douce J, Flament J, Escartin C, et al. New paradigm to assess brain cell morphology by diffusion-weighted MR spectroscopy in vivo. *Proceedings of the National Academy of Sciences* 2016;113(24):6671–6676.
- [46] Pena-Reyes CA, Sipper M. Evolutionary computation in medicine: an overview. *Artificial Intelligence in Medicine* 2000;19(1):1–23.

- [47] Poot DH, Arnold J, Achten E, Verhoye M, Sijbers J. Optimal experimental design for diffusion kurtosis imaging. *IEEE transactions on medical imaging* 2010;29(3):819–829.
- [48] Ronen I, Budde M, Ercan E, Annese J, Techawiboonwong A, Webb A. Microstructural organization of axons in the human corpus callosum quantified by diffusion-weighted magnetic resonance spectroscopy of N-acetylaspartate and post-mortem histology. *Brain Structure and Function* 2014;219(5):1773–1785.
- [49] Sabat S, Mir R, Guarini M, Guesalaga A, Irarrazaval P. Three dimensional k-space trajectory design using genetic algorithms. *Magnetic resonance imaging* 2003;21(7):755–764.
- [50] Santis S, Assaf Y, Evans CJ, Jones DK. Improved precision in CHARMED assessment of white matter through sampling scheme optimization and model parsimony testing. *Magnetic resonance in medicine* 2014;71(2):661–671.
- [51] Sen PN. Time-dependent diffusion coefficient as a probe of geometry. *Concepts in Magnetic Resonance Part A* 2004;23(1):1–21.
- [52] Sepehrband F, O'Brien K, Barth M. A time-efficient acquisition protocol for multipurpose diffusion-weighted microstructural imaging at 7 Tesla. *Magnetic Resonance in Medicine* 2017;.
- [53] Sotiropoulos SN, Jbabdi S, Xu J, Andersson JL, Moeller S, Auerbach EJ, et al. Advances in diffusion MRI acquisition and processing in the Human Connectome Project. *Neuroimage* 2013;80:125–143.
- [54] Spillman R. Solving large knapsack problems with a genetic algorithm. In: *Systems, Man and Cybernetics, 1995. Intelligent Systems for the 21st Century.*, IEEE International Conference on, vol. 1 IEEE; 1995. p. 632–637.
- [55] Stanisz GJ, Wright GA, Henkelman RM, Szafer A. An analytical model of restricted diffusion in bovine optic nerve. *Magnetic Resonance in Medicine* 1997;37(1):103–111.
- [56] Stejskal EO, Tanner JE. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *The journal of chemical physics* 1965;42(1):288–292.
- [57] Tuch DS. Q-ball imaging. *MR in medicine* 2004;52(6):1358–1372.
- [58] Wedeen VJ, Hagmann P, Tseng WYI, Reese TG, Weisskoff RM. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic resonance in medicine* 2005;54(6):1377–1386.
- [59] Wu YC, Field AS, Alexander AL. Computation of Diffusion Function Measures in q -Space Using Magnetic Resonance Hybrid Diffusion Imaging. *IEEE transactions on medical imaging* 2008;27(6):858–865.
- [60] Zhang H, Schneider T, Wheeler-Kingshott CA, Alexander DC. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. *Neuroimage* 2012;61(4):1000–1016.

[Figure 8 about here.]

[Figure 9 about here.]

LIST OF FIGURES

- 1 Exhaustive search results of the *optimization by shells* for the in silico experiment with $n_{max} = 100$. The plots at the top present all the 658,008 feasible acquisition schemes arranged from best to worst, illustrating the mean squared errors (MSEs) of signal reconstruction (top-left plot) and the normalized Hamming distances from the global optimum ± 1 standard deviation (top-right). In order to visualize the analyzed (G, Δ) parameter space, the percentiles $pc = 0\%, 1\%, 10\%, 50\%, 90\%$ are annotated on both plots, showing respectively the global optimum, the top 1% solutions, the top 10% solutions, etc. The corresponding cumulative averages of acquisition schemes are depicted in the heat maps at the bottom. The colors reflect the likelihood of a given (G, Δ) pair in the scheme. The heat maps for $pc \leq 0\%$ and $pc \leq 1\%$ represent, respectively, the global optimum and its proximity. The interval between $pc = 10\%$ and $pc = 90\%$ contains a huge spectrum of schemes with similar MSEs and almost equally large distances from the global optimum. 18
- 2 The convergence of our method is reached before the 30th iteration of the algorithm. The plots present the convergence curves averaged over 30 runs of SGA in the *optimization by shells* scenario for the in silico experiment with the Rician noise, SNR=20db (the upper plot), and the C57Bl6 wild-type mouse corpus callosum body (the lower plot). The time budgets $n_{max} \in \{100, 200, 300, 400\}$ are color-coded. . 19
- 3 Our approach significantly outperforms the other two subsampling schemes (with p-value $< 10^{-5}$) in both granularity levels: *optimization by shells* (top-row plots) and *optimization by measures* (bottom-row plots), reaching lower means and standard deviations. The plots present the residuals of the dMRI signal reconstruction for the in silico experiment with the dispersion controlled by the concentration parameter $\kappa = 4$ and SNR=20 (13dB) (left-column plots) and the body part of C57B16 wild-type mouse corpus callosum (right-column plots), for the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The results are expressed as normalized root mean squared errors (NRMSEs) of signal reconstruction with standard deviations (STDs) aggregated over 30 runs. 20
- 4 The concentrations of high likelihood (red squares) tend to form consistent shapes that are especially visible in the *optimization by shells* scenario. The plots present the averages of the top 10% acquisition schemes found by SGA in the *optimization by shells* scenario (top row) and the *optimization by measures* scenario (bottom row) for the in silico experiment with the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The colors reflect the likelihood of a given (Δ, G) pair in the scheme. 21
- 5 The concentrations of high likelihood (red squares) tend to form consistent shapes that are especially visible in the *optimization by shells* scenario. The plots present the averages of the top 10% acquisition schemes found by SGA in the *optimization by shells* scenario (top row) and the *optimization by measures* scenario (bottom row) for the *body* region of the C57Bl6 wild-type mouse corpus callosum, and the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The colors reflect the likelihood of a given (Δ, G) pair in the scheme. 22
- 6 All the tested schemes give comparable index curves for the in silico experiment. The plots present the reconstruction of spatio-temporal indices RTOP, RTAP, RTPP, MSD ± 1 standard deviation with $n_{max} = 100$. The black plots show reference curves obtained from the densely acquired signal, the color plots represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes. As expected, the RTOP, RTAP, and RTPP indices are decreasing, whereas MSD is increasing with diffusion time. 23

7 Our approach ensures the least dispersed outputs. The plots present the reconstruction of spatio-temporal indices RTOP, RTAP, RTPP, MSD ± 1 standard deviation with $n_{max} = 100$, obtained for the three regions of C57Bl6 wild-type mouse corpus callosum (CC). The black plots show reference curves obtained from the densely acquired signal, the color plots represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes. As expected, the RTOP, RTAP, and RTPP indices are decreasing, whereas MSD is increasing with diffusion time. 24

S1 Our approach ensures the least dispersed outputs. The plots present the reconstruction of Mean Diffusivity, Axial Diffusivity, Radial Diffusivity, and Fractional Anisotropy indices ± 1 standard deviation with $n_{max} = 100$ in the *optimization by measures*, obtained for the in silico experiment either with or without Rician noise. The black plots show reference curves obtained from the densely acquired signal, the color plots represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes. 25

S2 Our approach ensures the least dispersed outputs. The plots present the reconstruction of Mean Diffusivity, Axial Diffusivity, Radial Diffusivity, and Fractional Anisotropy indices ± 1 standard deviation with $n_{max} = 100$, obtained for the three regions of C57Bl6 wild-type mouse corpus callosum (CC). The black plots show reference curves obtained from the densely acquired signal, the color plots represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes. 26

Exhaustive search results

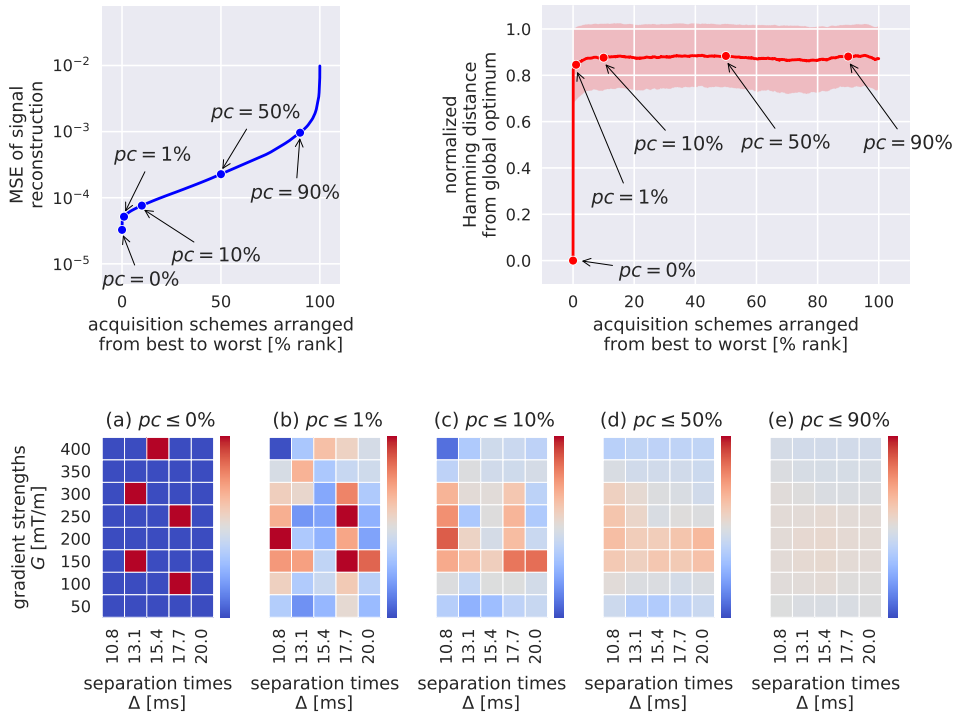


FIGURE 1 Exhaustive search results of the *optimization by shells* for the in silico experiment with $n_{max} = 100$. The plots at the top present all the 658,008 feasible acquisition schemes arranged from best to worst, illustrating the mean squared errors (MSEs) of signal reconstruction (top-left plot) and the normalized Hamming distances from the global optimum ± 1 standard deviation (top-right). In order to visualize the analyzed (G, Δ) parameter space, the percentiles $pc = 0\%$, 1% , 10% , 50% , 90% are annotated on both plots, showing respectively the global optimum, the top 1% solutions, the top 10% solutions, etc. The corresponding cumulative averages of acquisition schemes are depicted in the heat maps at the bottom. The colors reflect the likelihood of a given (G, Δ) pair in the scheme. The heat maps for $pc \leq 0\%$ and $pc \leq 1\%$ represent, respectively, the global optimum and its proximity. The interval between $pc = 10\%$ and $pc = 90\%$ contains a huge spectrum of schemes with similar MSEs and almost equally large distances from the global optimum.




FIGURE 2 The convergence of our method is reached before the 30th iteration of the algorithm. The plots present the convergence curves averaged over 30 runs of SGA in the *optimization by shells* scenario for the in silico experiment with the Rician noise, SNR=20db (the upper plot), and the C57Bl6 wild-type mouse corpus callosum body (the lower plot). The time budgets $n_{max} \in \{100, 200, 300, 400\}$ are color-coded.




FIGURE 3 Our approach significantly outperforms the other two subsampling schemes (with p -value $< 10^{-5}$) in both granularity levels: *optimization by shells* (top-row plots) and *optimization by measures* (bottom-row plots), reaching lower means and standard deviations. The plots present the residuals of the dMRI signal reconstruction for the in silico experiment with the dispersion controlled by the concentration parameter $\kappa = 4$ and SNR=20 (13dB) (left-column plots) and the body part of C57Bl6 wild-type mouse corpus callosum (right-column plots), for the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The results are expressed as normalized root mean squared errors (NRMSEs) of signal reconstruction with standard deviations (STDs) aggregated over 30 runs.

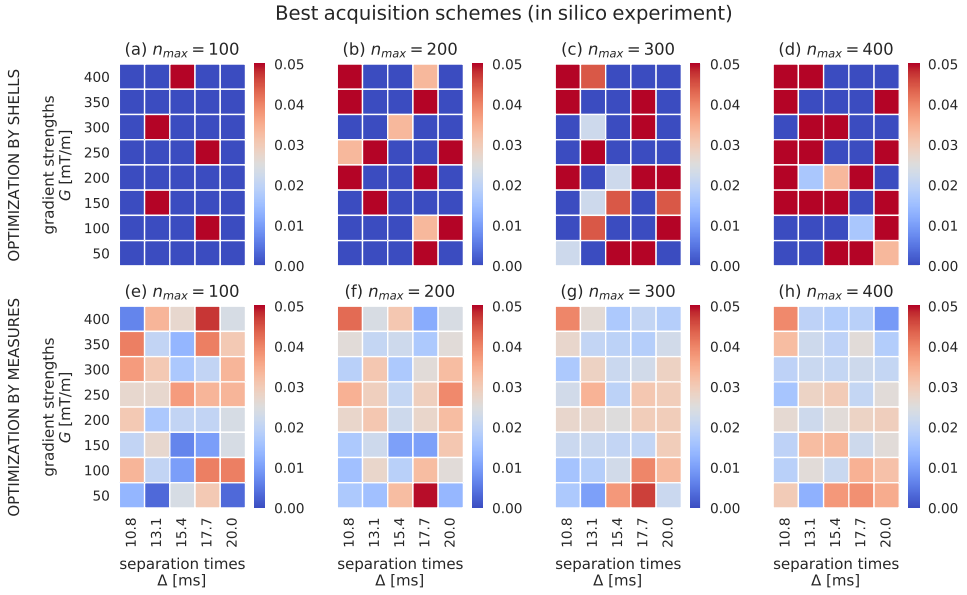


FIGURE 4 The concentrations of high likelihood (red squares) tend to form consistent shapes that are especially visible in the *optimization by shells* scenario. The plots present the averages of the top 10% acquisition schemes found by SGA in the *optimization by shells* scenario (top row) and the *optimization by measures* scenario (bottom row) for the in silico experiment with the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The colors reflect the likelihood of a given (Δ, G) pair in the scheme.




FIGURE 5 The concentrations of high likelihood (red squares) tend to form consistent shapes that are especially visible in the *optimization by shells* scenario. The plots present the averages of the top 10% acquisition schemes found by SGA in the *optimization by shells* scenario (top row) and the *optimization by measures* scenario (bottom row) for the *body* region of the C57Bl6 wild-type mouse corpus callosum, and the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The colors reflect the likelihood of a given (Δ, G) pair in the scheme.

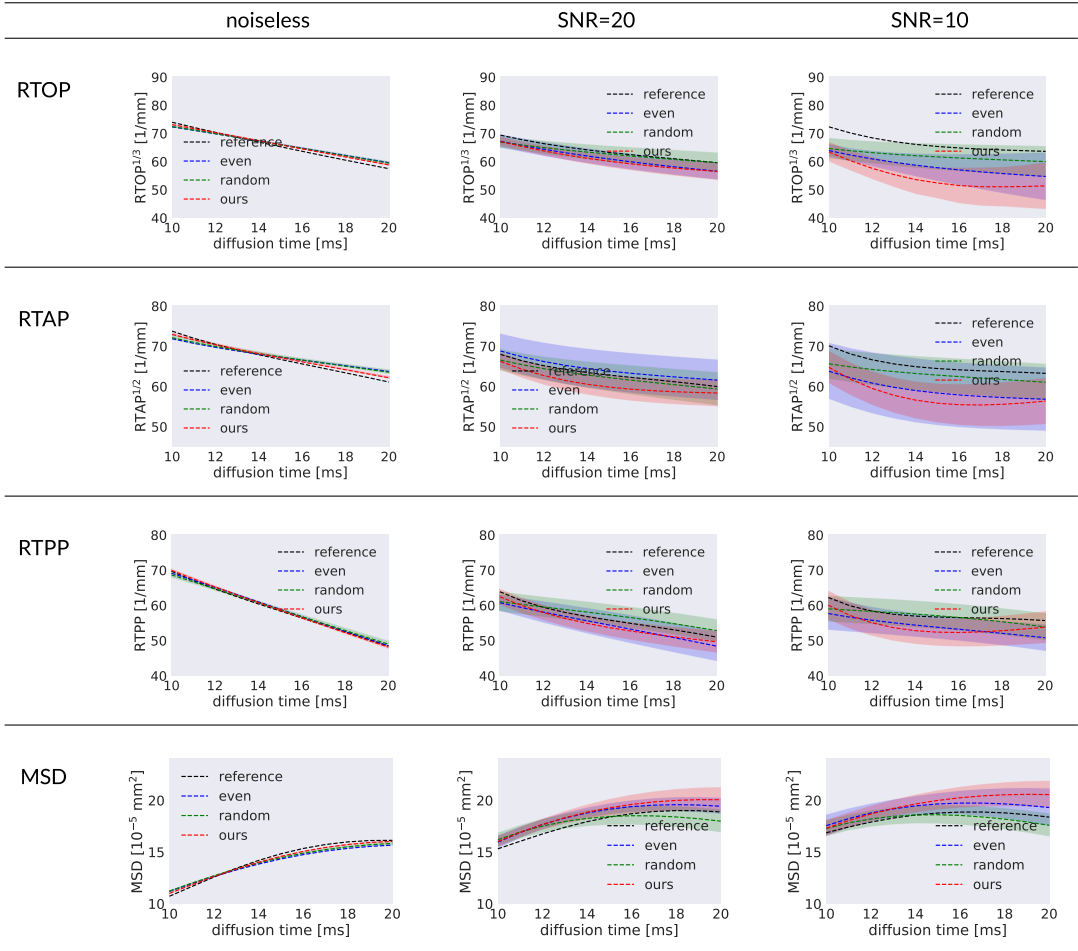


FIGURE 6 All the tested schemes give comparable index curves for the in silico experiment. The plots present the reconstruction of spatio-temporal indices RTOP, RTAP, RTPP, MSD ± 1 standard deviation with $n_{max} = 100$. The black plots show reference curves obtained from the densely acquired signal, the color plots represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes. As expected, the RTOP, RTAP, and RTPP indices are decreasing, whereas MSD is increasing with diffusion time.

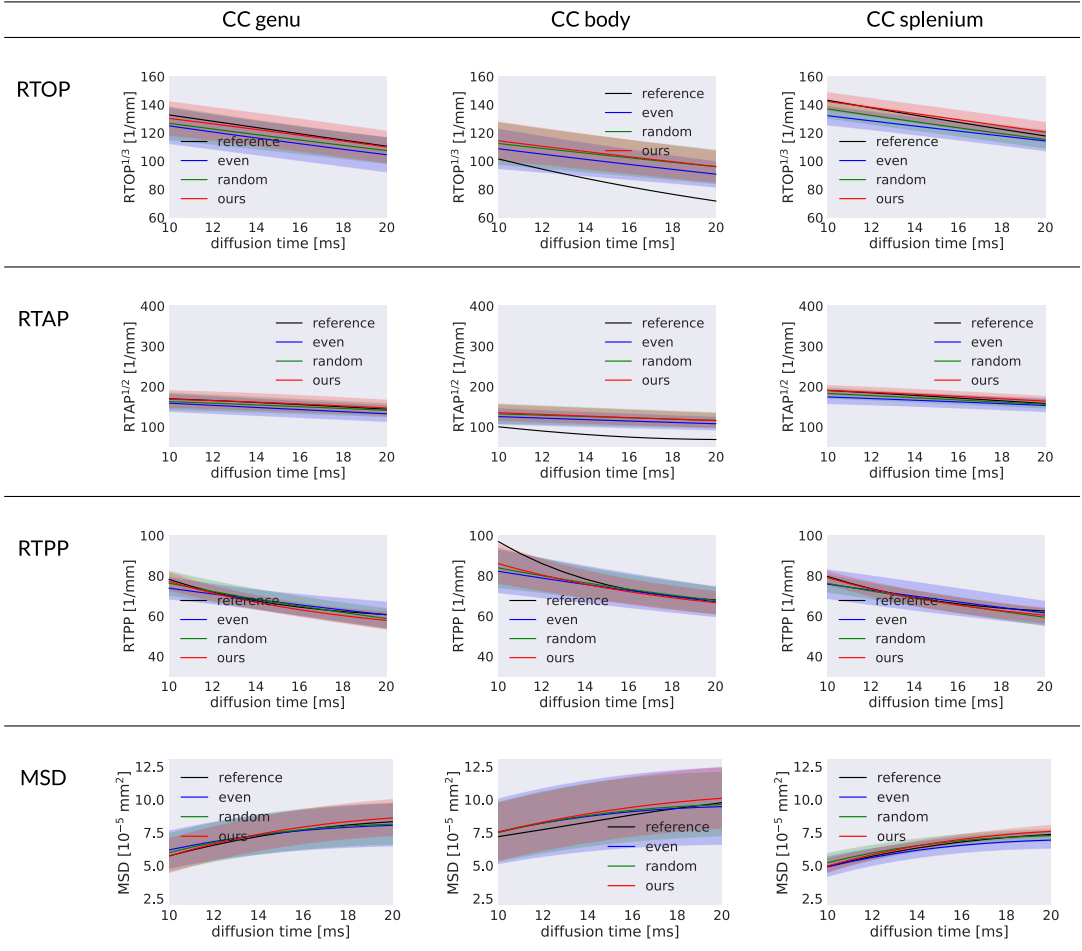


FIGURE 7 Our approach ensures the least dispersed outputs. The plots present the reconstruction of spatio-temporal indices RTOP, RTAP, RTPP, $\text{MSD} \pm 1$ standard deviation with $n_{max} = 100$, obtained for the three regions of C57Bl6 wild-type mouse corpus callosum (CC). The black plots show reference curves obtained from the densely acquired signal, the color plots represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes. As expected, the RTOP, RTAP, and RTPP indices are decreasing, whereas MSD is increasing with diffusion time.

SUPPORTING INFORMATION FIGURE S1 Our approach ensures the least dispersed outputs. The plots present the reconstruction of Mean Diffusivity, Axial Diffusivity, Radial Diffusivity, and Fractional Anisotropy indices ± 1 standard deviation with $n_{max} = 100$ in the *optimization by measures*, obtained for the in silico experiment either with or without Rician noise. The black plots show reference curves obtained from the densely acquired signal, the color plots represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes.

SUPPORTING INFORMATION FIGURE S2 Our approach ensures the least dispersed outputs. The plots present the reconstruction of Mean Diffusivity, Axial Diffusivity, Radial Diffusivity, and Fractional Anisotropy indices ± 1 standard deviation with $n_{max} = 100$, obtained for the three regions of C57Bl6 wild-type mouse corpus callosum (CC). The black plots show reference curves obtained from the densely acquired signal, the color plots represent *even* (blue), *random* (green), and *ours* (red) subsampling schemes.

LIST OF TABLES

- 1 Parameters used for generating the in silico diffusion data with the values taken from the referenced literature. Only the principal orientation was chosen arbitrary. The Watson distribution is an antipodally symmetric distribution, centered around the principal orientation μ , describing a density controlled by concentration parameter $\kappa > 0$, which is inversely related to the dispersion of cylinders. 28
- 2 Our approach significantly outperforms the other two subsampling schemes (with p-value $< 10^{-5}$) in both granularity levels (*by shells* and *by measures*), reaching lower means and standard deviations. The tables present the residuals of the dMRI signal reconstruction for the in silico experiment and the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The results are expressed as normalized root mean squared errors (NRMSEs) of signal reconstruction with standard deviations (STDs) aggregated over 30 runs. The STD is omitted in the *even* scheme which is deterministic and thus its STD=0. 29
- 3 Our approach significantly outperforms the other two subsampling schemes (with p-value $< 10^{-5}$) in both granularity levels (*by shells* and *by measures*), reaching lower means and standard deviations. The tables present the residuals of the dMRI signal reconstruction for the three regions of C57Bl6 wild-type mouse corpus callosum (CC) and the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The results are expressed as normalized root mean squared errors (NRMSEs) of signal reconstruction with standard deviations (STDs) aggregated over 30 runs. The STD is omitted in the *even* scheme which is deterministic and thus its STD=0. 30

intra-cellular fraction			extra-cellular fraction		
parameter's name	value	ref.	parameter's name	value	ref.
fraction weight	0.3	[11]	fraction weight	0.7	[11]
principal orientation μ	azim.: $\pi/2$, elev.: 0		principal orientation μ	azim.: $\pi/2$, elev.: 0	
parallel diffusivity λ_{\parallel}	$1.7 \times 10^{-9} \text{ m}^2/\text{s}$	[2]	parallel diffusivity λ_{\parallel}	$1.7 \times 10^{-9} \text{ m}^2/\text{s}$	[2]
cylinder diameter d	$1.0 \times 10^{-6} \text{ m}$	[2]	bulk diffusivity constant λ_{∞}	$6.5 \times 10^{-10} \text{ m}^2/\text{s}$	[11]
concentration parameter κ	4	[60]	characteristic coefficient A	$7.41 \times 10^{-12} \text{ m}^2$	[11]

TABLE 1 Parameters used for generating the in silico diffusion data with the values taken from the referenced literature. Only the principal orientation was chosen arbitrary. The Watson distribution is an antipodally symmetric distribution, centered around the principal orientation μ , describing a density controlled by concentration parameter $\kappa > 0$, which is inversely related to the dispersion of cylinders.

	time	NRMSE \pm STD [$\times 10^{-3}$]					
noise	budget	optimization by shells			optimization by measures		
	n_{max}	ours	random	even	ours	random	even
	100	17.3 \pm 0.72	43.7 \pm 22.70	23.5	14.9 \pm 0.14	23.3 \pm 3.23	28.1
noiseless	200	13.2 \pm 0.26	21.8 \pm 4.85	15.7	12.7 \pm 0.13	15.8 \pm 0.71	21.9
signal	300	12.1 \pm 0.15	18.5 \pm 4.35	16.4	12.2 \pm 0.07	14.5 \pm 0.53	14.0
	400	11.8 \pm 0.06	16.2 \pm 2.55	13.1	11.7 \pm 0.08	14.0 \pm 0.46	14.3
	100	165.6 \pm 1.06	225.4 \pm 40.15	192.1	166.6 \pm 0.41	184.4 \pm 2.49	186.6
SNR=20	200	161.5 \pm 0.39	196.8 \pm 29.45	190.4	162.7 \pm 0.21	177.9 \pm 3.03	183.8
(13dB)	300	158.6 \pm 0.15	184.3 \pm 18.75	171.1	160.5 \pm 0.18	173.2 \pm 1.75	172.4
	400	157.3 \pm 0.18	174.5 \pm 9.16	169.6	159.0 \pm 0.22	170.2 \pm 2.12	171.9
	100	275.6 \pm 1.51	329.0 \pm 28.08	304.5	273.6 \pm 0.64	298.5 \pm 4.64	304.7
SNR=10	200	266.8 \pm 0.48	301.1 \pm 20.03	299.5	268.4 \pm 0.43	288.8 \pm 3.41	297.8
(10dB)	300	263.7 \pm 0.28	291.2 \pm 16.55	281.3	266.5 \pm 0.97	283.6 \pm 3.47	283.0
	400	262.1 \pm 0.51	283.7 \pm 13.34	279.7	264.6 \pm 0.18	279.4 \pm 2.24	280.3

TABLE 2 Our approach significantly outperforms the other two subsampling schemes (with p-value $< 10^{-5}$) in both granularity levels (*by shells* and *by measures*), reaching lower means and standard deviations. The tables present the residuals of the dMRI signal reconstruction for the in silico experiment and the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The results are expressed as normalized root mean squared errors (NRMSEs) of signal reconstruction with standard deviations (STDs) aggregated over 30 runs. The STD is omitted in the *even* scheme which is deterministic and thus its STD=0.

region	time	NRMSE \pm STD [$\times 10^{-3}$]					
of	budget	optimization by shells			optimization by measures		
interest	n_{max}	ours	random	even	ours	random	even
	100	110.2 \pm 0.45	131.6 \pm 14.98	120.1	106.9 \pm 0.33	118.9 \pm 1.95	120.8
CC	200	103.7 \pm 0.17	112.7 \pm 4.18	112.1	102.8 \pm 0.14	109.8 \pm 1.10	114.1
genu	300	101.7 \pm 0.14	108.3 \pm 3.11	106.9	101.5 \pm 0.09	106.4 \pm 0.78	106.5
	400	100.6 \pm 0.05	104.9 \pm 1.75	105.4	100.8 \pm 0.13	104.1 \pm 0.52	104.1
	100	127.6 \pm 0.85	172.2 \pm 44.30	144.2	122.9 \pm 0.54	142.1 \pm 3.28	155.3
CC	200	116.1 \pm 0.45	136.2 \pm 9.83	132.7	116.6 \pm 0.19	128.4 \pm 1.76	149.7
body	300	111.9 \pm 0.36	124.8 \pm 6.41	122.0	113.4 \pm 0.20	121.4 \pm 1.73	124.1
	400	109.9 \pm 0.22	118.9 \pm 3.13	119.1	111.5 \pm 0.10	116.8 \pm 1.16	119.9
	100	107.0 \pm 1.09	133.3 \pm 19.75	123.1	104.4 \pm 0.37	117.9 \pm 2.80	131.1
CC	200	101.4 \pm 0.41	111.2 \pm 4.82	109.6	100.9 \pm 0.11	108.4 \pm 1.32	141.0
splenium	300	99.9 \pm 0.42	105.7 \pm 2.23	106.8	99.6 \pm 0.09	104.2 \pm 0.75	104.4
	400	99.0 \pm 0.42	103.2 \pm 1.46	103.7	98.6 \pm 0.12	101.9 \pm 0.53	103.8

TABLE 3 Our approach significantly outperforms the other two subsampling schemes (with p-value $< 10^{-5}$) in both granularity levels (*by shells* and *by measures*), reaching lower means and standard deviations. The tables present the residuals of the dMRI signal reconstruction for the three regions of C57Bl6 wild-type mouse corpus callosum (CC) and the time budgets $n_{max} \in \{100, 200, 300, 400\}$. The results are expressed as normalized root mean squared errors (NRMSEs) of signal reconstruction with standard deviations (STDs) aggregated over 30 runs. The STD is omitted in the *even* scheme which is deterministic and thus its STD=0.