

COARSE-GRAINED SPATIOTEMPORAL ACQUISITION DESIGN FOR DIFFUSION MRI

P. Filipiak¹, R. Fick¹, A. Petiet², M. Santin², A.-C. Philippe², S. Lehericy², R. Deriche¹, D. Wassermann³

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

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

³Inria, CEA, Université Paris-Saclay, France

ABSTRACT

Acquisition protocols that allow to capture time-dependent changes in diffusion signal require long imaging time. We address this issue through an optimized subsampling scheme that maximizes accuracy of the spatiotemporal diffusion signal representation, $q\tau$ -dMRI, for given time constraints. Our proposed coarse-grained variant of the problem reduces the space of feasible acquisition parameters compared to the fine-grained approach causing no significant deterioration of a reconstruction accuracy in most of the studied cases.

1. INTRODUCTION

We address the problem of spatiotemporal acquisition design for diffusion Magnetic Resonance Imaging (dMRI) with the goal to reduce imaging time. Most of the protocols that allow to capture time-dependent changes in diffusion signal require long acquisition [1], which hampers their applicability in the clinical practice.

In this vein, Alexander [2] proposed an optimized acquisition scheme that maximized accuracy of brain white matter microstructure reconstruction for a given time limit by tuning axon radii estimation. Caruyer et al. [3] suggested a dMRI acquisition scheme with the uniform coverage of gradient directions in order to assure rotation invariance of a sampled signal. Recently, we proposed an optimized acquisition design that takes, for the first time, into account time-dependence in dMRI [4] by applying the spatiotemporal signal representation, called $q\tau$ -dMRI, introduced by Fick et al. [5].

In this work, we propose the *coarse-grained* variant of the spatiotemporal dMRI acquisition design. This change of granularity lets us reduce the size of feasible acquisition parameter space compared to our previous *fine-grained* variant, which simplifies finding optimal subsampling schemes.

Our experiments present time-dependent diffusion signals fitted with $q\tau$ -dMRI on crossing fibers in both synthetic and in vivo data of the C57Bl6 wild-type mouse.

This work has received funding from the ANR/NSF award NeuroRef; the MAXIMS grant funded by ICM's The Big Brain Theory Program, ANR-10-IAIHU-06, and the European Research Council (ERC) under the Horizon 2020 research and innovation program (ERC Advanced Grant agreement No 694665: CoBCoM).

2. THEORY

In dMRI, the Ensemble Average Propagator (EAP) $P(\mathbf{R}; \tau)$ holds 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 Diffusion Weighted Images (DWIs), which are obtained by applying two sensitizing diffusion gradients $\mathbf{G} \in \mathbb{R}^3$ of pulse length $\delta > 0$, separated by the time interval $\Delta > 0$. Assuming narrow pulses ($\delta \rightarrow 0$), the EAP is related to the dMRI signal through the Inverse Fourier Transform defined as

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

with the signal attenuation $E(\mathbf{q}, \tau)$ measured at the diffusion encoding position $\mathbf{q} = \gamma\delta\mathbf{G}/2\pi$ and the diffusion time $\tau = \Delta - \delta/3$, where γ is the gyromagnetic ratio.

We reconstruct the continuous EAP from a finite set of DWIs by using the spatiotemporal signal representation $q\tau$ -dMRI [5]. For this, we estimate the functional basis coefficients $\mathbf{c} \in \mathbb{R}^{N_q \times N_\tau}$, where N_q and N_τ are the maximum expansion orders of the spatial and temporal bases, respectively. Let $\hat{E}(\mathbf{q}, \tau, \mathbf{c})$ be the $q\tau$ -dMRI representation of $E(\mathbf{q}, \tau)$. We find \mathbf{c} by applying the following GraphNet regularization

$$\begin{aligned} \arg \min_{\mathbf{c}} & \iint [E(\mathbf{q}, \tau) - \hat{E}(\mathbf{q}, \tau, \mathbf{c})]^2 d\mathbf{q}d\tau + \\ & + \lambda \iint [\nabla^2 \hat{E}(\mathbf{q}, \tau, \mathbf{c})]^2 d\mathbf{q}d\tau + \alpha \|\mathbf{c}\|_1 \end{aligned} \quad (2)$$

subject to $\hat{E}(0, \tau, \mathbf{c}) = 1$ and $\hat{E}(\mathbf{q}, 0, \mathbf{c}) = 1$,

where $\lambda, \alpha > 0$ stand for the smoothness and sparsity regularization weights, respectively.

3. METHODS

Assume we have a dense dMRI acquisition protocol containing $K > 0$ shells, each of which comprises of $S_k > 0$ sampling points ($k = 1, \dots, K$). For simplicity of notation, we enumerate all the samples from 1 to $N = \sum_{k=1, \dots, K} S_k$. Our aim is to find a subsampling scheme composed of (not more than) $n_{max} > 0$ samples among the N ones used in the dense acquisition protocol, such that the loss of reconstructed $q\tau$ -dMRI signal representation is minimized.

3.1. Fine-grained problem [4]

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 is included in the subsampling scheme, whereas $x_i = 0$ determines its exclusion. Formally, our goal is to solve the following optimization problem

$$\begin{aligned} \arg \min_x F(x) &= \frac{1}{M} \sum_{j=1}^M \|E^{(j)} - \hat{E}_x^{(j)}\|_2^2 \\ \text{subject to } \sum_{i=1}^N x_i &\leq n_{max} \quad \text{with } 1 \leq n_{max} \leq N - 1, \end{aligned} \quad (3)$$

where $M > 0$ is the number of voxels in each DWI, $E^{(j)}$ is the normalized signal attenuation captured in the j -th voxel, and $\hat{E}_x^{(j)}$ is the corresponding signal obtained by fitting the $q\tau$ -dMRI representation (defined in Equation 2) 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.

As we showed before [4], the problem posed in Equation (3) is NP-hard, which generally precludes finding exact global optima. In this paper, we propose to coarse this problem in such a way that the obtained subsampling schemes are rotation invariant and less subject-specific.

3.2. Coarse-grained problem

Recall the K -shell dense acquisition introduced earlier. For all $k = 1, \dots, K$, let us define a constant $y_k = x_{S_{k-1}+1} + \dots + x_{S_k}$ (with $S_0 = 0$) holding a number of samples taken into the subsampling scheme from the k -th shell. In other words, for a given scheme x we aggregate the samples as follows

$$\underbrace{(x_1, \dots, x_{S_1})}_{y_1}, \underbrace{(x_{S_1+1}, \dots, x_{S_1+S_2})}_{y_2}, \dots, \underbrace{(x_{S_1+\dots+S_{K-1}+1}, \dots, x_N)}_{y_K}$$

and thus obtain a vector $y = (y_1, \dots, y_K) \in (\mathbb{Z}_+ \cup \{0\})^K$.

Let us now consider the inverse scenario: For a given y we construct a subsampling scheme $x'(y) \in \{0, 1\}^N$ such that $\sum_{i=1}^N x'_i(y) \leq n_{max}$ and for all $k = 1, \dots, K$ the coefficient y_k determines the number of samples from the k -th shell in x' . The mapping $y \mapsto x'(y)$ postulated above is not unique, so we impose additionally a uniform distribution of samples in order to assure rotation invariance of the subsampling scheme, as suggested by Alexander [2]. Such uniformity can be defined in numerous ways [3], e.g. as uniform per shell or among all shells combined. In this paper, we assume the *uniform per shell* case, i.e. for each $k = 1, \dots, K$ we want the y_k samples to be spread uniformly across the k -th shell.

Mathematically, our coarse-grained problem assumes finding a vector y such that $x'(y)$ satisfies the conditions given in Equation (3). By introducing this modification, we shrink the parameter space from $\binom{N}{n_{max}}$ to maximally $\binom{K+n_{max}-1}{K-1}$ feasible schemes.




Fig. 1. Regions of interest studied in the in vivo experiment.

3.3. Experiments

We used Population Based Incremental Learning (PBIL) [7] to solve both optimization problems: fine- and coarse-grained. PBIL is an evolutionary algorithm that was shown effective in constrained high-dimensional binary optimization problems like ours [8]. In most of the studied cases, the convergence was reached after 20-30 iterations with the population size 100 and the remaining parameters set to their default values. We chose experimentally the regularization weights $\lambda = 10^{-5}$ and $\alpha = 0.02$.

Our dense pre-acquisition covered $K = 40$ shells, each of which comprised of $S = 20$ directions and one b_0 -image, i.e. $N = 40 \times 20 = 800$ 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] 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/mm². 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. [3].

In silico data: As first data set, we generated time-dependent diffusion data using our Python-based [9] implementation of the two-compartment model with the intra-cellular fraction modeled with a set of Watson dispersed crossing cylinders [10] and the extra-cellular one modeled with a temporal zeppelin [1]. The angles of crossings were controlled by a random variable with a distribution centered at $\pi/2$ and a dispersion $\pi/4$. Following the study of Zhang et al. [10], we used the concentration parameters of the Watson distribution $\kappa \in \{0.25, 1, 4, 16\}$, where $\kappa = 16$ means highly concentrated and $\kappa = 0.25$ means highly dispersed cylinders. Additionally, we considered two variants of our data set perturbed with Rician noise having Signal-to-Noise Ratio (SNR) 20 and 10, respectively.

In vivo data: In the second data set, we used in vivo DWIs of the C57Bl6 wild-type mouse. All animal experi-

dispersion	SNR	$\kappa = 16$			$\kappa = 4$			$\kappa = 1$			$\kappa = 0.25$		
		∞	20	10	∞	20	10	∞	20	10	∞	20	10
$\kappa = 16$	∞	0.421	0.015	0.013	0.212	0.044	0.051	0.195	0.003	0.179	0.389	0.059	0.118
	20	0.195	0.165	0.950	0.122	0.119	0.333	0.341	0.408	0.444	0.352	0.361	0.620
	10	0.043	0.082	0.247	0.040	0.560	0.038	0.030	0.062	0.833	0.062	0.983	0.308
$\kappa = 4$	∞	0.139	0.001	0.027	0.111	0.020	0.010	0.345	0.001	0.008	0.002	0.006	0.212
	20	0.028	0.080	0.074	0.077	0.219	0.454	0.081	0.702	0.976	0.014	0.169	0.651
	10	0.001	0.319	0.629	0.001	0.845	0.552	0.001	0.850	0.281	0.001	0.534	0.535
$\kappa = 1$	∞	0.344	0.094	0.108	0.081	0.104	0.006	0.419	0.001	0.026	0.011	0.027	0.121
	20	0.010	0.146	0.489	0.017	0.529	0.033	0.016	0.071	0.700	0.013	0.406	0.217
	10	0.005	0.999	0.909	0.005	0.521	0.154	0.065	0.247	0.104	0.027	0.364	0.175
$\kappa = 0.25$	∞	0.636	0.406	0.299	0.942	0.660	0.586	0.301	0.309	0.611	0.301	0.275	0.997
	20	0.003	0.265	0.015	0.002	0.652	0.015	0.008	0.120	0.871	0.002	0.152	0.123
	10	0.014	0.533	0.303	0.007	0.658	0.914	0.003	0.932	0.424	0.006	0.446	0.307

Table 1. The p-values of the relative Student t-test for comparing coarse- and fine-grained problems using the mean squared residuals obtained in 30 runs of our optimization scheme on each of them for $n_{max} = 100$. The rows present (κ, SNR) -pairs of the in silico training data, while the columns stand for the testing data sets. The cases where the fine-grained variant outperforms the coarse-grained one under the confidence level $1 - \alpha = 0.99$ are printed in bold. All of them appear in noiseless cases only.

ments 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*.

For each of the 12 in silico (κ, SNR) -pairs we generated





Fig. 2. Our coarse-grained problem largely reduces the size of parameter space compared to the fine-grained one.



(a) in silico ($\kappa = 1$, noiseless) (b) in vivo (test ROI #1)

Fig. 3. Mean squared residuals averaged over 30 runs. Both optimized schemes give almost identical results and outperform even subsampling in each case.

training and testing data comprising of $M = 10$ voxels each. For the in vivo experiment, we chose 4 Regions of Interest (ROIs) presented in Figure 1. The training data ROI was located in the cerebrum area with crossings of projections and association fibers. The testing data ROIs #1-#3 were located in the thalamus, brain stem, and cerebellum (respectively).

For the acquisition time constraints, we considered 8 budget sizes $n_{max} \in \{20, 40, 60, 80, 100, 200, 300, 400\}$ out of 800 DWIs. The remaining 40 b_0 -images were excluded from the optimization process and used intact.

4. RESULTS & DISCUSSION

Our main goal is to find an acquisition design that maximizes reconstruction accuracy of the $q\tau$ -dMRI signal representation under given time constraints. We define our task as the optimization problem that we solve on two granularity levels.

Figure 2 illustrates the sizes of acquisition parameter space in the fine- and coarse-grained variants of the problem. Our proposed coarse-grained approach largely reduces the search space, which simplifies finding optimal subsampling scheme defined in Equation 3. For instance, we reduce the set of feasible solutions from 3.41×10^{129} to 5.05×10^{34} for $n_{max} = 100$. As we show later, this reduction causes no significant deterioration in accuracy of $q\tau$ -dMRI signal representation, except for the cases with noiseless data.

Figure 3 shows representative examples of mean squared residuals of $q\tau$ -dMRI signal representation averaged over 30 runs, obtained under various time budgets n_{max} on both in silico and in vivo data. In the latter case, we use the densely acquired data as reference. For comparison, we present the results of *even subsampling*, i.e. taking every second sample from the dense scheme for $n_{max} = N/2$, every fourth for $n_{max} = N/4$, etc. As Figure 3 shows, both our *optimized* schemes, namely *coarse* and *fine*, give almost identical reconstruction accuracy in all cases and each time outperform *even subsampling*. Also note that for $n_{max} \geq 200$ the residuals

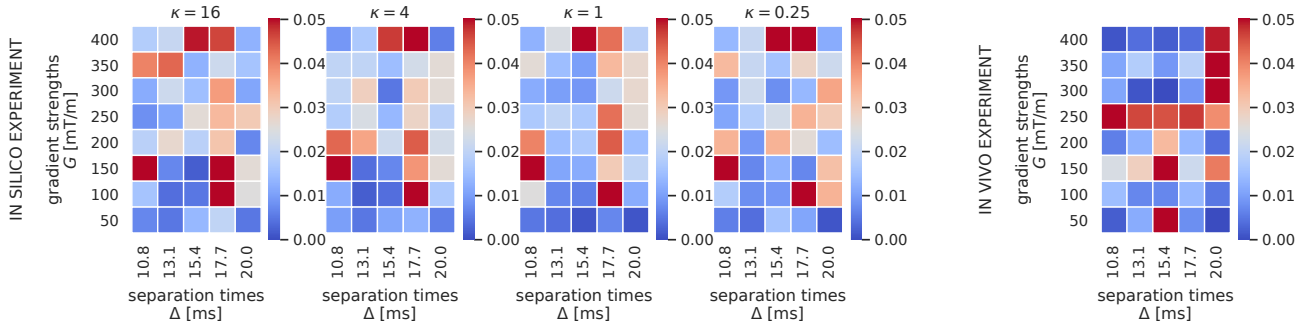


Fig. 4. Averages of 10 best acquisition schemes found with our coarse-grained optimization for $n_{max} = 100$ on noiseless in silico data (plots on the left) with $\kappa = 16, 4, 1, 0.25$, respectively, and in vivo testing data ROI #1 (plot on the right). The distribution of most frequently used parameter pairs (G, Δ) in our optimized subsampling schemes form clusters of red and orange squares concentrated in the same 3-4 areas of the parameter space.

of optimized schemes remain nearly constant, which suggests that the optimization is most profitable for small time budgets.

Table 1 summarizes p-values of the relative Student t-test that we used for comparing coarse- and fine-grained problems with the mean squared residuals obtained in 30 runs of our optimization scheme on each of them, for $n_{max} = 100$. The rows present (κ, SNR) -pairs of the in silico training data, while the columns stand for the testing data sets. The cases where the fine-grained variant outperforms the coarse-grained one under the confidence level $1 - \alpha = 0.99$ are printed in bold. Note that in most of the cases there are no significant differences between respective results. The fine-grained optimization is superior to the coarse-grained one only when trained and/or tested on certain noiseless data sets. Furthermore, the p-values in analogous comparison for in vivo data were $(0.017, 0.081, 0.026)$ for the testing data ROIs #1-#3, respectively, which means there were no significant differences between fine- and coarse-grained variants.

Figure 4 illustrates averages of 10 best acquisition schemes found with our coarse-grained optimization for $n_{max} = 100$ on noiseless in silico data (plots on the left) with $\kappa = 16, 4, 1, 0.25$, respectively, and in vivo testing data ROI #1 (plot on the right). The distribution of most frequently used parameter pairs (G, Δ) in our optimized subsampling schemes form clusters of red and orange squares concentrated in the same 3-4 areas of the parameter space. This shows that our coarse-grained optimization converges repeatedly to the same optima promising stability of identified subsampling schemes.

5. CONCLUSION

In this work, we proposed the *coarse-grained* variant of the spatiotemporal dMRI acquisition design. Our scheme was able to reduce acquisition time even 8 times ($n_{max} = N/8 = 100$) with minimized accuracy loss when representing crossing fibers with q_T -dMRI. Moreover, our course-grained optimization allowed to simplify subsampling schemes in such a way that they produced only the numbers of gradi-

ent directions per shell, assuming uniform distribution of sampling points on the hemisphere, rather than specifying their exact coordinates. This helps make acquisition schemes less subject-specific and thus applicable immediately in other studies. Our experiments revealed the limitations of the coarse-grained approach on the noiseless synthetic data, although in practice dMRI signal is relatively noisy and as such gives comparable results in both granularity levels of our optimization.

6. REFERENCES

- [1] L.M. Burcaw, E. Fieremans, and D.S. Novikov, "Mesoscopic structure of neuronal tracts from time-dependent diffusion," *NeuroImage*, 2015.
- [2] D.C. Alexander, "A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features," *Magnetic Resonance in Medicine*, 2008.
- [3] E. Caruyer, C. Lenglet, G. Sapiro, and R. Deriche, "Design of multishell sampling schemes with uniform coverage in diffusion MRI," *Magnetic Resonance in Medicine*, 2013.
- [4] P. Filippiak, R. Fick, A. Petiet, M. Santin, A.-C. Philippe, S. Lehericy, P. Ciuciu, R. Deriche, and D. Wassermann, "Reducing the number of samples in spatio-temporal dMRI acquisition design," *MRM* (in press).
- [5] R. Fick, A. Petiet, M. Santin, A.-C. Philippe, S. Lehericy, R. Deriche, and D. Wassermann, "Non-parametric GraphNet-regularized representation of dMRI in space and time," *Medical Image Analysis*, 2018.
- [6] H. Gudbjartsson and S. Patz, "The Rician distribution of noisy MRI data," *Magnetic Resonance in Medicine*, 1995.
- [7] S. Baluja and R. Caruana, "Removing the genetics from the standard genetic algorithm," *Machine Learning*, 1995.
- [8] R. Sagarna and P. Larrañaga, "Solving the 0-1 knapsack problem with EDAs," *Estimation of Distribution Algorithms*, 2002.
- [9] R. Fick, D. Wassermann, and R. Deriche, "Mipy: An open-source framework to improve reproducibility in brain microstructure imaging," *OHBM 2018-Human Brain Mapping*, 2018.
- [10] H. Zhang, T. Schneider, C.A. Wheeler-Kingshott, and D.C. Alexander, "Noddi: practical in vivo neurite orientation dispersion and density imaging of the human brain," *NeuroImage*, 2012.