



An Evolutionary Framework for Microstructure-Sensitive Generalized Diffusion Gradient Waveforms

Truffet, Raphaël; Rafael-Patino, Jonathan; Girard, Gabriel; Pizzolato, Marco; Barillot, Christian; Thiran, Jean Philippe; Caruyer, Emmanuel

Published in:
Medical Image Computing and Computer Assisted Intervention

Link to article, DOI:
[10.1007/978-3-030-59713-9_10](https://doi.org/10.1007/978-3-030-59713-9_10)

Publication date:
2020

Document Version
Early version, also known as pre-print

[Link back to DTU Orbit](#)

Citation (APA):
Truffet, R., Rafael-Patino, J., Girard, G., Pizzolato, M., Barillot, C., Thiran, J. P., & Caruyer, E. (2020). An Evolutionary Framework for Microstructure-Sensitive Generalized Diffusion Gradient Waveforms. In A. L. Martel, P. Abolmaesumi, D. Stoyanov, D. Mateus, M. A. Zuluaga, S. K. Zhou, D. Racoceanu, & L. Joskowicz (Eds.), *Medical Image Computing and Computer Assisted Intervention* (pp. 94-103). Springer.
https://doi.org/10.1007/978-3-030-59713-9_10

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

An evolutionary framework for microstructure-sensitive generalized diffusion gradient waveforms

Raphaël Truffet^{1*}, Jonathan Rafael-Patino^{2*}, Gabriel Girard^{2,3,4}, Marco Pizzolato^{2,5,6}, Christian Barillot¹, Jean-Philippe Thiran^{2,3,4}, and Emmanuel Caruyer¹

¹ Univ Rennes, Inria, CNRS, Inserm, IRISA UMR 6074, Empenn ERL U-1228, F-35000, Rennes, France

² Signal Processing Lab (LTS5), École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

³ Radiology Department, Centre Hospitalier Universitaire Vaudois (CHUV) and University of Lausanne (UNIL), Lausanne, Switzerland

⁴ Center for BioMedical Imaging (CIBM), Lausanne, Switzerland

⁵ Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kongens Lyngby, Denmark

⁶ Danish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University, Hospital Hvidovre, Hvidovre, Denmark

Abstract. In diffusion-weighted MRI, general gradient waveforms became of interest for their sensitivity to microstructure features of the brain white matter. However, the design of such waveforms remains an open problem. In this work, we propose a framework for generalized gradient waveform design with optimized sensitivity to selected microstructure features. In particular, we present a rotation-invariant method based on a genetic algorithm to maximize the sensitivity of the signal to the intra-axonal volume fraction. The sensitivity is evaluated by computing a score based on the Fisher information matrix from Monte-Carlo simulations, which offer greater flexibility and realism than conventional analytical models. As proof of concept, we show that the optimized waveforms have higher scores than the conventional pulsed-field gradients experiments. Finally, the proposed framework can be generalized to optimize the waveforms for to any microstructure feature of interest.

Keywords: diffusion MRI · acquisition design · generalized gradient waveforms · Fisher information · Monte-Carlo simulations.

1 Introduction

By measuring the displacement of molecules at the micrometer scale, diffusion-weighted MRI (DW-MRI) is sensitive to the fine structure of biological tissues.

* These two authors contributed equally

Biophysical models have been proposed to extract useful information about tissue microstructure *in vivo*. In particular in the brain white matter, axons have a coherent organization. Besides estimating the orientation of fibers [21], one can quantify the apparent intra-axonal volume fraction (IAVF), parameters of the distribution of axon radii [5, 6] or orientation coherence of the fiber orientations [22].

Complementary to analytical modeling of the DW-MRI signal, Monte-Carlo simulation offers increased flexibility for the description of biological substrates and higher accuracy of the DW-MRI signal for selected microstructure configurations [15, 14, 18]. The signals computed using Monte-Carlo simulations were shown to be a good predictor of those measured *in vivo* [19].

Another key factor, besides modeling, to improve the accuracy of microstructure features estimation is the experimental design. In [4], the authors present a general framework for experimental design, applied to the optimization of pulsed gradient spin echo (PGSE) sequence parameters for the estimation of the composite hindered and restricted model of diffusion (CHARMED). The same framework was later extended to oscillating gradient spin echo (OGSE) [12], generalized gradient waveforms (linear-encoding) [11] and gradient trajectories (spherical-encoding) [10].

In this work, we focus on the choice of gradient waveforms to optimize the sensitivity of the acquisition to the intra-axonal volume fraction (IAVF) in white matter. Using a three-dimensional substrate model of white matter and Monte-Carlo simulations [18], we optimize the gradient waveforms for an increased rotation-averaged Fisher information. The search for optimized waveforms follows a particle swarm heuristic. Our optimized waveforms are compared with the PGSE sequence.

2 Methods

In this section, we present our novel framework for the generation of waveforms sensitive to IAVF, beginning by designing a score measuring this sensitivity. We then describe the numerical substrate, and how they can be used to compute the scores. Finally, we describe the evolution approach that leads to optimized waveforms.

2.1 A score based on the Fisher information

The Fisher information is a measure of how sensitive a gradient trajectory, \mathbf{g} , is to a microstructure parameter of interest. Indeed, the Cramér-Rao bound, defined as the lower bound on the variance of any unbiased estimator, is computed as the inverse of the Fisher information [13]. In this study, we focus on the IAVF, f , while the other microstructure parameters are fixed (see Section 2.2). We focused on the IAVF for two main reasons: i- the IAVF is a biologically relevant index of axonal loss, and therefore a biomarker specific to several neurodegenerative diseases; ii- measuring the IAVF from clinically plausible datasets is more

accessible and less sensitive to noise than other parameters of interest, such as e.g. axon diameter index, for which a precise estimation remains challenging for diameters below 5 μm . We derive in the sequel the Fisher information for f , and provide implementation details for its computation using Monte-Carlo simulations.

Signal and noise model The signal attenuation model $A(f, \mathbf{g})$ depends on a the IAVF, f , and an effective DW-MRI gradient trajectory, $\mathbf{g}(t), t \in [0, TE]$ (TE is the echo time). Besides, the magnitude signal in DW-MRI is corrupted by noise. As mentioned in [3], it is important to consider Rician noise rather than Gaussian noise, since the latter leads to an unrealistic choice of higher b-value. In some situations, the Rician model does not adequately describe noise properties. However, in most cases, it can be considered a valid model [2]. The probability density function of the measured signal, \tilde{A} , in noise is, with a spread parameter σ :

$$p(\tilde{A}; f, \mathbf{g}, \sigma^2) = \frac{\tilde{A}}{\sigma^2} I_0 \left(\frac{A(f, \mathbf{g}) \tilde{A}}{\sigma^2} \right) \exp \left(-\frac{A^2(f, \mathbf{g}) + \tilde{A}^2}{2\sigma^2} \right). \quad (1)$$

Computation of the Fisher information The Fisher information for the parameter f is defined as

$$F_0(\mathbf{g}, f, \sigma^2) = E \left[\frac{\partial^2 \log p}{\partial f^2}(\tilde{A}; f, \mathbf{g}, \sigma^2) \right]. \quad (2)$$

The full derivation for the Fisher information F_0 (not reported here) can be found in the Appendix of [3]; it can be computed from the estimation of the partial derivatives $\partial A / \partial f$. Note that in our case, there is no analytical expression for the signal attenuation; instead, A is computed using the Monte-Carlo simulator. Because of the intrinsic uncertainty in the signal, we cannot estimate these partial derivatives using classical finite-difference. Therefore, we empirically propose to perform linear regression from a set of signals generated for 10 values of f around the value of interest.

We note that the Fisher information depends on the value of the IAVF f . Without prior knowledge on this parameter, and to avoid introducing any bias in the acquisition design, we compute the average Fisher information over a set $\Omega = \{0.25, 0.50, 0.75\}$:

$$F(\mathbf{g}, \sigma^2) = \frac{1}{|\Omega|} \sum_{f \in \Omega} F_0(\mathbf{g}, f, \sigma^2). \quad (3)$$

Gradient waveforms and rotation invariance In this work, we focus on gradient trajectories with a fixed orientation, $\mathbf{u} \in \mathcal{S}^2$:

$$\mathbf{g}(t) = g(t) \mathbf{u}.$$

We want to separate the search for optimal waveforms from the search for optimal sampling directions, the latter having already received prior attention in the community [16, 8, 9, 7]. Besides, to have a rotation-invariant measure, we define the sampling score, U , as the average over directions on the sphere:

$$U(g, \sigma^2) = \int_{S^2} F(g \mathbf{u}, \sigma^2) d^2 \mathbf{u}. \quad (4)$$

In practice, the integral in Eq. 4 is computed by taking advantage of the cylindrical symmetry of the substrates, and considering a set of nine gradient directions making an angle θ with the cylinders uniformly spread in the range $[0, \pi/2]$. As a result, we needed to run 270 simulations (3 IAVF values, each with 10 values for the linear regression and each with 9 angles θ) to compute the score $U(g, \sigma^2)$ for one waveform. Overall, this corresponds to 30 Monte Carlo particle trajectories files to be generated.

2.2 Numerical substrate design

All substrates were generated assuming a two-compartment model composed of intra- and extra-axonal spaces. The intra-axonal space compartment was represented using a collection of parallel cylinders with radius sampled from a Gamma distribution $\Gamma(2.0, 0.35)$ (in μm), which is in the range of values reported from histology samples [1, 17]. The extra-axonal space compartment corresponds to space outside the cylindrical axons, which comprise the extra-axonal matrix, glial cells, and cerebrospinal fluid. The substrates were then generated by randomly placing a total of 10^4 sampled cylinders into an isotropic voxel, without intersection between them, and ensuring periodicity at the voxel boundaries [18]. Fig. 1, panel a), shows a toy example of a numerical substrate.

A total of 40 numerical substrates were generated for four IAVF of 0.25, 0.5, 0.6, and 0.75. For each of the four nominal IAVF, we generated 10 samples around the target value to compute the numerical derivative of the Fisher information. To generate samples with the same intra-axonal configuration, but slightly different IAVF, we first generated one randomly packed phantom for each of the four selected IAVF, and then scaled each one of them by a small factor $1.0 + \epsilon$, with $\epsilon = \{0.00, 0.001, 0.002, \dots, 0.010\}$, without scaling the cylinders radii distribution, as is shown in Fig. 1, panel b). This way, the same intra-axonal space configuration is kept, but different IAVF can be computed. Doing so, we avoid any bias arising from radii re-sampling and positioning of the cylinders, which would effectively change the considered geometry, hence the signal.

Principle To build waveforms with the highest score defined in Eq. 4, we perform a stochastic optimization based on a genetic algorithm. To narrow down the search space, the admissible range of b -values is pre-determined. Intuitively, if the b -value is too low, there is almost no diffusion-weighting; conversely, a too high b -value leads to a highly attenuated signal which has an amplitude comparable to that of the noise floor. We start with an initial set of 100 random waveforms, and we perform cross-overs to build the next generation.




Fig. 1: Example of substrate design composed of intra- and extra- axonal spaces. The intra-axonal space compartment is represented using a collection of parallel cylinders with various radii packed inside a containing voxel, which results in the top view displayed in panel a). Panel b) shows how the base configuration is then transformed by scaling the voxel and the position of the cylinders by a constant factor, but without scaling the cylinder’s radii; this transformation reduces the volume ratio between the voxel and the cylinders total volume, thus reducing the represented IAVF.

Initialization The process to build random gradient waveforms is based on Markov Chains: the value of the waveforms at time $t + dt$ only depends on the value of the waveform at time t . This method allows us to build waveforms that respect the properties of maximum gradient strength ($T \cdot m^{-1}$) and slew rate ($T \cdot m^{-1} \cdot s^{-1}$) that are well documented. Other limitations such as the duty cycle could have been added [20], but are not documented enough to add numeric constraints. We referred to the limitations of a SIEMENS Prisma 3T, for possible future experimentations:

$$|g(t)| \leq 80 \text{ mT} \cdot \text{m}^{-1} \quad ; \quad \left| \frac{dg}{dt}(t) \right| \leq 200 \text{ T} \cdot \text{m}^{-1} \cdot \text{s}^{-1}. \quad (5)$$

We generate the waveform between $t = 0$ and $t = TE/2$, and we perform a symmetry to obtain the full gradient trajectory. Examples of randomly generated gradients are shown on Fig. 2 (a). We generated several waveforms and selected only those that have a b -value in the targeted range.

Evolution process We defined the cross-over as the combination of two waveforms. We keep the beginning of a waveform, and the end of another, as shown in Fig. 2. The position of the cut is sampled randomly from a Gaussian centered on $TE/4$, with a standard deviation of $TE/12$ to avoid that the new waveform inherits too much from only one of the two previous waveforms.




Fig. 2: Illustration of the cross-over between two waveforms (a), g_1 and g_2 . We prepare two amplitude modulations (b) which will keep respectively the beginning and the end of waveforms g_1 and g_2 . After applying this modulation, we sum the results to obtain a new waveform (d). NB: we only represent the chronogram for $t \in [0, TE/2]$, *i.e.* before the 180° pulse.

2.3 Genetic evolution

The generation $i + 1$ is built on generation i . First, we randomly select two waveforms. Then, the probability to select a waveform is a linear scaling of the score that leads to have a probability to select the waveform with the highest score 10 times higher than selecting the waveform with the lowest score. If the b -value of the resulting waveform is outside the targeted range, we drop this cross-over and try another one. We repeat this process until having 95 waveforms that fall within the target range. The last 5 waveforms needed to complete generation $i + 1$ are randomly generated such as for generation 0 to promote novelty.

2.4 Signal simulation

The simulated DW-MRI signals were computed using DW-MRI Monte-Carlo simulations as described in [18]. However, since in each iteration of our proposed framework it is required to recompute the DW-MRI signal of all the numerical substrates at each generation of waveforms, this requires generating a high number of new signals for the same substrate. Because of this, we proposed to first generate and store all the Monte-Carlo particles dynamics required to compute the DW-MRI signal using the MC/DC open-source simulator [18], and then use a tailored version of this simulator written in C++ CUDA, which is able to compute the DW-MRI signal in a fraction of the original computation time. Using the latter approach, we were able to produce one generation of waveforms in about 30 minutes, in contrast to 20 hours for the former.

For each numerical substrate, we generated and stored a total of 2×10^5 particles, with a maximum diffusion time of 100 ms, and in time-steps intervals of 20 μ s. The diffusion coefficient was set to $D = 1.7 \times 10^{-9} \text{ m}^2 \cdot \text{s}^{-1}$ [3].

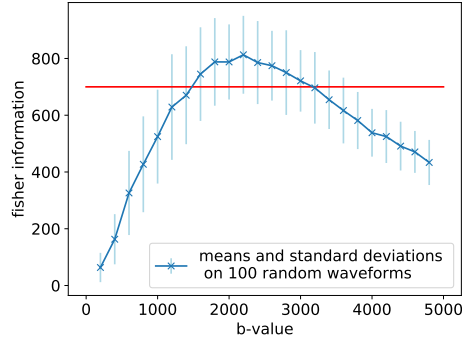


Fig. 3: Fisher information computed for a family of randomly generated gradient waveforms, spanning a range of b -values from 0 to $5,000 \text{ s} \cdot \text{mm}^{-2}$. The Fisher information shows a dependence on the b -values, but within a given b -value, there is a remaining variability which is explained by waveform. We restricted our waveform search in the b -value range of $[1500, 3100] \text{ s} \cdot \text{mm}^{-2}$ (value above the red line) to further investigate their sensitivity to the IAVF.

3 Results and Discussion

3.1 Pre-optimization of the b -value

In Fig. 3, we report the Fisher information for several b -values ranging from 0 to $5,000 \text{ s} \cdot \text{mm}^{-2}$, with 100 different waveforms for each b -value. We first notice a high variability within one b -value. This shows that the b -value is not enough to characterize the efficiency of one waveform for the estimation of the IAVF. Then, we also notice that the b -value still represents an important parameter since we observe, as expected, a low score for low b -values and too high b -values. This result leads us to restrict our search for waveforms to those within the target range $[1,500 \text{ s} \cdot \text{mm}^{-2}; 3,100 \text{ s} \cdot \text{mm}^{-2}]$.

3.2 Waveforms optimization

As we can see in Fig. 5a, after 30 generations, the score of the waveforms has increased compared to generation 0. Since most of the waveforms of generation 30 have a higher score than the random waveforms of generation 0, one can hope that this increase is due to the optimization process and not only to new random waveforms incorporated at each generation. The waveforms with the highest score of generation 30 are shown in Fig. 4.




Fig. 4: Waveforms genetically generated to optimize the Fisher Information. From left to right, the figure shows the 4 waveforms of generation 30 with the highest score.

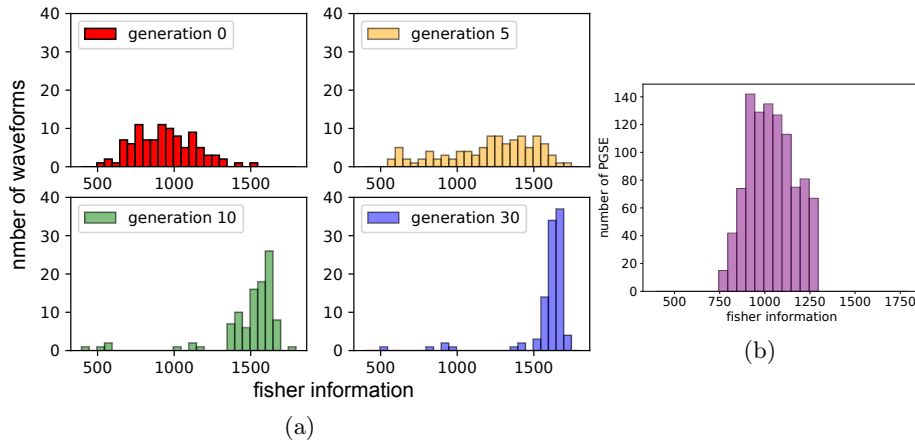


Fig. 5: Histogram of the Fisher information of DW-MRI sequences. (a) Histogram associated to gradient waveforms across four generations of the genetic algorithm. (b) Histogram computed for the family of 1,000 PGSE.

3.3 Comparison of score between PGSE and genetically optimized waveforms

As a proof-of-concept, we compared our optimized waveforms with PGSE. We computed the score for a family PGSE covering a wide range of timing and gradient amplitude parameters. We generated all the PGSE sequences with $G_{\max} = 80 \text{ mT} \cdot \text{m}^{-1}$; δ between 0.5 ms and 40 ms with a step of 0.5 ms; and Δ between 10 ms and 80 ms with a step of 0.5 ms. After filtering only feasible sequences, and only selecting those with a b -value between 1500 and 3100, we obtained approximately 1,000 PGSE. Fig. 5b shows the distribution of scores among these PGSE sequences.

The maximum score with these b -values is lower than the score obtained with generalized waveforms in generation 30. We can also notice that, even in generation 0, some randomly generated waveforms had a higher score than PGSE sequences. This shows that generalized waveforms have a greater potential than PGSE sequences at identifying the substrate’s IAVF.

4 Conclusion

The optimization of gradient waveforms is an important step in designing new acquisition sequences in DW-MRI. The generalized waveforms protocol obtained in this work is optimized independently from the sampling directions. We designed a genetic algorithm to make the waveforms evolve. After a few generations, our algorithm created a family of waveforms that have a higher sensitivity to the IAVF than PGSE.

We have presented a novel framework with several advantages. First, it is based on Monte-Carlo simulations, and thus can be adapted to a variety of microstructure configurations, such as substrate with various compartment shapes and sizes. Second, generalized waveform optimization can be performed on one or several microstructure parameters of interest. In this work, we optimized for the IAVF parameter in two-compartment substrates composed of parallel cylinders with various radii and packing densities, but other biomarkers can be used instead. The only requirement is our ability to compute partial derivatives of the signal with respect to these parameters. Future work will target angular dispersion and spherical pore size parameter estimation. In vivo experiments will be a necessary step for the validation of the generated waveforms. Designing an experiment that shows the efficiency of the waveforms without having the ground truth microstructure is a challenging issue.

Acknowledgements

This research project is part of the MMINCARAV Inria associate team program between Empenn (Inria Rennes Bretagne Atlantique) and LTS5 (École Polytechnique Fédérale de Lausanne — EPFL) started in 2019. Raphaël Truffet’s PhD is partly funded by ENS Rennes. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan XP GPU used for this research. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 754462.

References

1. Aboitiz, F., Scheibel, A.B., Fisher, R.S., Zaidel, E.: Fiber composition of the human corpus callosum. *Brain Research* **598**(1), 143 – 153 (1992). [https://doi.org/https://doi.org/10.1016/0006-8993\(92\)90178-C](https://doi.org/https://doi.org/10.1016/0006-8993(92)90178-C), <http://www.sciencedirect.com/science/article/pii/000689939290178C>
2. Aja-Fernandez, S., Vegas-Sanchez-Ferrero, G.: *Statistical Analysis of Noise in MRI: Modeling, Filtering and Estimation*. Springer Publishing Company, Incorporated, 1st edn. (2016)
3. Alexander, D.C.: A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magnetic Resonance in Medicine* **60**(2), 439–448 (aug 2008). <https://doi.org/10.1002/mrm.21646>, <http://doi.wiley.com/10.1002/mrm.21646>
4. Alexander, D.C.: A general framework for experiment design in diffusion mri and its application in measuring direct tissue-microstructure features. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **60**(2), 439–448 (2008)
5. Assaf, Y., Basser, P.J.: Composite hindered and restricted model of diffusion (charmed) mr imaging of the human brain. *Neuroimage* **27**(1), 48–58 (2005)
6. Assaf, Y., Blumenfeld-Katzir, T., Yovel, Y., Basser, P.J.: Axcaliber: a method for measuring axon diameter distribution from diffusion mri. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **59**(6), 1347–1354 (2008)
7. Bates, A.P., Khalid, Z., Kennedy, R.A.: An optimal dimensionality sampling scheme on the sphere with accurate and efficient spherical harmonic transform for diffusion mri. *IEEE Signal Processing Letters* **23**(1), 15–19 (2015)
8. Caruyer, E., Lenglet, C., Sapiro, G., Deriche, R.: Design of multishell sampling schemes with uniform coverage in diffusion mri. *Magnetic resonance in medicine* **69**(6), 1534–1540 (2013)
9. Cheng, J., Shen, D., Yap, P.T.: Designing single-and multiple-shell sampling schemes for diffusion mri using spherical code. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 281–288. Springer (2014)
10. Drobnjak, I., Alexander, D.C.: Optimising time-varying gradient orientation for microstructure sensitivity in diffusion-weighted mr. *Journal of Magnetic Resonance* **212**(2), 344–354 (2011)
11. Drobnjak, I., Siow, B., Alexander, D.C.: Optimizing gradient waveforms for microstructure sensitivity in diffusion-weighted mr. *Journal of Magnetic Resonance* **206**(1), 41–51 (2010)
12. Drobnjak, I., Zhang, H., Ianuș, A., Kaden, E., Alexander, D.C.: Pgse, ogse, and sensitivity to axon diameter in diffusion mri: insight from a simulation study. *Magnetic resonance in medicine* **75**(2), 688–700 (2016)
13. Galdos, J.I.: A Cramer-Rao Bound for Multidimensional Discrete-Time Dynamical Systems. *IEEE Transactions on Automatic Control* **25**(1), 117–119 (1980). <https://doi.org/10.1109/TAC.1980.1102211>
14. Ginsburger, K., Matuschke, F., Poupon, F., Mangin, J.F., Axer, M., Poupon, C.: Medusa: A gpu-based tool to create realistic phantoms of the brain microstructure using tiny spheres. *NeuroImage* **193**, 10–24 (2019)
15. Hall, M.G., Alexander, D.C.: Convergence and parameter choice for monte-carlo simulations of diffusion mri. *IEEE transactions on medical imaging* **28**(9), 1354–1364 (2009)

16. Jones, D.K., Horsfield, M.A., Simmons, A.: Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **42**(3), 515–525 (1999)
17. Lamantia, A.S., Rakic, P.: Cytological and quantitative characteristics of four cerebral commissures in the rhesus monkey. *Journal of Comparative Neurology* **291**(4), 520–537 (1990). <https://doi.org/10.1002/cne.902910404>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/cne.902910404>
18. Rafael-Patino, J., Romascano, D., Ramirez-Manzanares, A., Canales-Rodríguez, E.J., Girard, G., Thiran, J.P.: Robust Monte-Carlo Simulations in Diffusion-MRI: Effect of the Substrate Complexity and Parameter Choice on the Reproducibility of Results. *Frontiers in Neuroinformatics* **14**, 8 (mar 2020). <https://doi.org/10.3389/fninf.2020.00008>
19. Rensonnet, G., Scherrer, B., Girard, G., Jankovski, A., Warfield, S.K., Macq, B., Thiran, J.P., Taquet, M.: Towards microstructure fingerprinting: Estimation of tissue properties from a dictionary of monte carlo diffusion mri simulations. *NeuroImage* **184**, 964–980 (2019)
20. Sjölund, J., Szczepankiewicz, F., Nilsson, M., Topgaard, D., Westin, C.F., Knutsson, H.: Constrained optimization of gradient waveforms for generalized diffusion encoding. *Journal of Magnetic Resonance* **261**, 157–168 (2015). <https://doi.org/10.1016/j.jmr.2015.10.012>
21. Tuch, D.S.: Q-ball imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **52**(6), 1358–1372 (2004)
22. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: Noddi: practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* **61**(4), 1000–1016 (2012)