Diffusion MRI analysis

DOI:
10.26481/dis.20191025rh

Summary and Discussion

The primary aim of this thesis was to improve the robustness and efficiency of the methods used in diffusion MRI microstructure modeling. A secondary aim was to provide the community with software to utilize the findings of this work. For the primary aim we studied maximum likelihood estimation in Chapter 2, MCMC sampling in Chapter 3 and uncertainty estimation in Chapter 4. For the secondary aim we created the Microstructure Diffusion Toolbox, discussed in Chapter 5. The first section of this chapter summarizes the scientific contributions of each of the individual chapters. The second section discusses this thesis and provides an outlook for future work in microstructure modeling.

6.1 Summary

Chapter 2: Robust and fast nonlinear optimization of diffusion MRI microstructure model

Most of the authors in microstructure modeling prefer the use of non-linear optimization routines for their microstructure modeling. Often mentioned optimization routines are the Levenberg-Marquardt (Jespersen et al., 2010; Assaf, Freidlin, et al., 2004; Assaf, Blumenfeld-Katzir, et al., 2008) and Nelder-Mead Simplex routine (Landman, Bazin, and Prince, 2007). The initialization strategy is often not specified, although some authors mention a grid search prior to optimization (H. Zhang, Schneider, et al., 2012). In Chapter 2 we studied these two aspects of non-linear optimization, the choice of optimization routine and the choice of initialization strategy.

We tested three optimization routines against each other on a multitude of models and datasets, the Levenberg-Marquardt, Nelder-Mead simplex and a newly proposed routine, the Powell method. Based on robustness, accuracy, precision and run time we recommend the Powell method as preferred optimization routine for dMRI modeling. We recommend using a single optimization strategy for all microstructure models in a study to minimize possible bias from the optimization strategy.

We recommend the use of cascades as an initialization strategy. In cascaded optimization, the parameters of a complex model are initialized with corresponding and already estimated parameters of a simpler model. Applied recursively, this can lead to a cascade of models to fit, with each new model slightly more complex than the previous model. As an extension to cascaded initialization, we proposed fixing the diffusion orientation estimates of a model using orientations from Ball&Stick estimates. Fixing the orientation parameters reduces the degrees of freedom in a model, improving run time and fitting robustness of the remaining parameters.

Chapter 3: Robust and fast Markov Chain Monte Carlo sampling of diffusion MRI microstructure models

Some authors prefer the use of MCMC sampling for their model fitting purposes (Behrens et al., 2003; Sotiropoulos et al., 2013; Wegmann, Eklund, and Villani, 2017). Whereas maximum likelihood estimation provides only a point estimate of the fitted model parameters, MCMC recovers the entire posterior distribution of the model parameters given the data. In addition to a point estimate, the posterior distribution provides parameter uncertainties and parameter correlations. The aim of chapter 3 was to provide a generally applicable, scientifically validated MCMC methodology for microstructure modeling. To this end, we studied several methodological aspects of MCMC including the choice of proposal distribution, burn-in, thinning and number of samples to draw.

There are many different MCMC strategies proposed in the literature, each defined uniquely by their proposal distribution. In Chapter 3 we focused on self adjusting univariate proposal distributions and recommend the Adaptive Metropolis Within Gibbs routine based on its theoretical soundness and its general robust performance. For burn-in and thinning we provide empirical and theoretical arguments against their use and recommend to use them sparingly, if at all. To determine the number of MCMC samples to draw, we propose the effective sample size theory to predict the number of samples required for the estimate of the posterior to reach a predetermined precision.

Chapter 4: Fast quantification of uncertainty in non-linear diffusion MRI models

Parameter uncertainties are typically ignored in diffusion MRI studies, although recognized for their benefits in functional MRI (Chen et al., 2012; Woolrich et al., 2004). In Chapter 4 we discuss and compare two methods for computing parameter uncertainties, the Fisher Information Matrix (FIM) and MCMC sampling. We then show how these uncertainty estimates could be used for artifact detection and increasing the power in group studies.

We first compared the FIM against gold standard MCMC using multiple datasets and models. We show that the parameter uncertainties obtained using the FIM show close correspondence to that of MCMC, yet computed in only a fraction of the time. We then proposed two use cases for the computed parameter uncertainties, artifact detection and variance reduction in group statistics. Results show that parameter uncertainties are often high in regions with white matter acquisition artifacts, allowing the uncertainties to be used for single subject artifact detection. For group studies we recommend the use of variance weighted averaging. Doing so can reduce the overall variance in group statistics and reduce the effect of data artifacts without discarding data from the analysis. Since the power of a study depends partially on the variance in the estimates, lowering the variance in group statistics increases the power of a study.

Chapter 5: MDT, the Microstructure Diffusion Toolbox

To make the findings in this thesis reproducible and easily applicable to other studies, we additionally present our research findings in the form of a software package. This software package, the Microstructure Diffusion Toolbox (MDT) is dedicated to diffusion MRI microstructure modeling. Since the models and algorithms in MDT have been discussed in the previous chapters, Chapter 5 focuses on the design concepts, user interfaces and application examples of MDT.

This software package contains all the optimization and MCMC sampling routines mentioned in this thesis. By default, it uses the Powell optimization routine and the Adaptive Metropolis Within Gibbs sampling routine for all microstructure fitting and MCMC sampling purposes. In addition, MDT automatically computes parameter uncertainties using the Fisher Information Matrix after every microstructure model optimization.

Unique to MDT is the combination of high level model definitions with high performance parallel processing. For maximum performance, we re-implemented all models and methods for execution on graphical processing units (GPU's). Using GPU's, MDT can process thousands of voxels simultaneously, making it one of the fastest microstructure modeling tools currently available. To hide away most of the GPU programming complexity, we created a microstructure modeling framework. Users can interact with this modeling framework using the high level Python programming language, while MDT creates the necessary GPU code in the background