Diffusion MRI analysis

Valorisation

Knowledge valorisation refers to the "process of creating value from knowledge, by making knowledge suitable and/or available for social (and/or economic) use and by making knowledge suitable for translation into competitive products, services, processes and new commercial activities" - as detailed in "*Regulation governing the attainment of doctoral degrees*", Maastricht University (2018). As the work presented in this thesis can potentially be applied in commercial products, in this chapter I take the opportunity to elaborate on the various valorisable aspects of the work in this thesis.

Diffusion MRI (dMRI) is one of the preferred approaches for investigating the brain's white matter microstructure in vivo. In a scientific setting, dMRI can be used to further a structural understanding of the human brain, while in a clinical setting it can be used as a diagnostic instrument for detecting, for instance, acute ischemia (lack of oxygen) in the human brain.

Due to the nature of a dMRI scan, the measured signal provides only an indirect view of the underlying cellular structures. This makes modeling essential in order to extract quantitative measures from diffusion MRI data. Many models have been proposed over the years, each differing in underlying assumptions and applicable tissue types. To infer information from the dMRI signal, most of these models need to be fitted to the data. This model fitting typically depends on non-linear analysis methods which have their respective problems in robustness, accuracy, precision and run-time. Improving on these can have a large impact in both the scientific and clinical domain.

The choice of algorithms

Since the quality and run-time of the dMRI analysis depend on the methods used, the investigation in the choice of algorithms can save both time and money.

In the first two chapters of this thesis we compared various optimization and Markov Chain Monte Carlo (MCMC) sampling algorithms on their usefulness for dMRI modeling. We noted that a smart choice of algorithm can heavily influence the run-time and quality of the results. For example, by showing that burn-in and thinning in MCMC sampling should be avoided, we can save hours of computation time for every dataset. In addition, we showed that using the same optimization routine for different models can improve the quality of the model comparisons. Since the used optimization routines influences the results, using the same routine for different models removes a possible confound from a

study, increasing the effect size and leading to a lower requirement in required participants.

The results from the third chapter also allow for a potential reduction of required participants in a dMRI study. The use of weighted averaging can reduce the effect of white matter artifacts in the averaged results. As such, subjects that previously would have been removed from consideration can now still be included in the group statistics as the weighted averaging will take care of removing or dampening the effect of the white matter artifacts. In addition, weighted averaging promises the lowest possible standard deviation of the mean. This in turn provides for a higher effect size when comparing two different populations. Both these effects can reduce the number of participants required for a study, reducing the number of expensive dMRI scans required.

Advantages of GPU computing

Next to the choice of analysis methods, the specific implementation can affect the analysis computation times. Since every increase in (d)MRI resolution leads to a power of three increment in the number of voxels, we need faster hardware and software to keep up with future datasets. Recently, Graphical Processing Units (GPU's) have seen increased usage in data sciences due to their large parallel computation possibilities. The use of a GPU is however not straightforward as software needs to be specifically adapted to take advantages of their compute power.

To take advantage of graphical processing units in dMRI modeling, we developed multiple open-source software packages with native support for GPU computations. The first of these packages, the MOT package, contains GPU accelerated implementations of all the non-linear optimization routines, all Markov Chain Monte Carlo sampling methods and the Fisher Information Matrix computation as used in this thesis. The second software package, the MDT package, contains highly optimized implementations of a large set of dMRI models. By making these routines open-source we bring all the advantages of this thesis to the community and wider public.

The reduction in computation time associated with GPU computing has several advantages. First of all, it allows researchers to analyze datasets faster, allowing for a higher throughput in modeling analysis. Second, it allows for more and faster iterations when developing new dMRI models, making research more time efficient. As a third advantage, it could reduce the need for a compute cluster. Since GPU's are between 30 and 60 times more efficient then a central processing unit (CPU) for dMRI modeling (see Chapter 4), providing researchers with a

single good GPU can prevent the need for a CPU cluster, potentially saving large amounts of money.

The modeling framework

Although GPU computing is a large selling point of the MDT and MOT software packages presented in this work, the MDT package has another useful feature, a large scale modeling framework. This modeling framework provides researchers with several advantages. Foremost, by bundling various dMRI models in one software package we allow for model comparisons using a single optimization routine for all models. This will increase effect size, leading to a reduction of required participants. Another useful aspect is that by providing a model building framework we allow researchers to reuse existing model implementations. This can save time when developing new models as the existing models are highly optimized and considered correct.

As a final selling point, the presented software packages can perhaps form the bridge between the scientific and the clinical domains. If the optimization and MCMC sampling routines can be verified for use in a medical setting, the MDT and MOT software packages could be used a model implementation framework to bring new dMRI models to the clinic.

Beyond diffusion MRI

One of the software packages written for this thesis is a stand-alone non-linear optimization toolbox with parallel processing capabilities, the Multithreaded Optimization Toolbox (MOT). While this software was essential to accelerating the computations in diffusion MRI analysis, it is in essence a general purpose optimization toolbox, capable of being extended to various other domains of science. For example, already within the field of magnetic resonance imaging, one may think of accelerating functional MRI computations, quantitative magnetic transfer models and structural MRI models. Outside of MRI, MOT may prove of use in genetic data analysis, protein folding, cancer cell research, microscopy analysis or any other field where a large amount of parallelizable optimization computations are required.