Glioma Microstructure Modeling from Diffusion MRI: A Self-Supervised Deep Learning Approach

Thea Våge

Master's Thesis
MSc Medical Technology
Department of Physics and Technology
University of Bergen

June, 2023

Abstract

Purpose: Gliomas are a highly heterogeneous group of primary brain tumors with poor prognosis, and treatment monitoring is challenging with its current diagnostic tool being invasive biopsy. In recent years, diffusion-weighted magnetic resonance imaging (dMRI) has become a widely used non-invasive technique that infers tissue microstructure in tumors. Such analysis presents opportunities in cancer diagnosis, tumor grading, and monitoring treatment effectiveness. A technique called vascular, extracellular and restricted diffusion for cytometry in tumors (VERDICT) has been demonstrated to be feasible for inferring glioma microstructure, but current methods are time-consuming in terms of acquisition and analysis and therefore not clinically applicable. In this thesis, the purpose was therefore to investigate whether VERDICT microstructure fitting in glioma tissue could be approached with a newly introduced self-supervised fully connected neural network method.

Methods: A large-scale synthetic dataset was simulated using the VERDICT signal equation based on known ground-truth glioma tissue parameters, with a highly detailed acquisition protocol. A feed-forward neural network (FFNN) was constructed to predict four tissue parameters from dMRI signals: the cell radius, intracellular (IC) volume fraction, extracellular-extravascular (EES) volume fraction and vascular volume fraction. The network was trained and tuned on the large-scale simulated dataset before it was trained on clinically applicable sequences and tested on *in vivo* dMRI. Furthermore, experiments were conducted to test the network's performance with different acquisition protocols, aiming to determine the potential improvement that could be achieved through protocol optimization. Lastly, the proposed ANN was compared to existing methods in terms of computational time and accuracy.

Results: The model trained on the large-scale synthetic data managed to estimate the four parameters with high precision, and the prediction time in one voxel was less than $10^{-4}s$ when applying a trained network to an *in vivo* dMRI. Predictions on acquired *in*

vivo data were compatible with tissue parameters from the literature for a longer dMRI scan, but the model showed lower accuracy in terms of fitting data with a shorter acquisition (i.e. less rich diffusion protocol). From experimenting with different acquisition protocols, the best performance was found on the longest protocol, but it was found that the accuracy of predictions was also dependent on the individual acquisition parameter ranges.

Conclusion: This project successfully demonstrated the ability to use the FFNN approach to fit the VERDICT model to dMRI data. Findings indicated that the FFNN fitting is heavily influenced by the acquisition protocol, and the method is not suitable for MRI acquisitions with less than 20 measurements. However, the method showed good potential for use in acquisitions with 436 measurements, and for predicting a radius of up to 10 μm, 145 high b-value measurements can be sufficient. Optimization of the acquisition schemes suggested that alternative schemes could further enhance the effectiveness of the technique. The FFNN showed several advantages in terms of computational time and accuracy of predictions compared to existing methods.

Acknowledgements

I would like to thank my supervisors Frank Riemer and Renate Grüner for granting me access to this project, and for the opportunity to be a part of the inspiring research environment at Mohn Medical Imaging and Visualization Center (MMIV). Thank you for sharing your knowledge with me and inspiring me with new ideas for the project. Especially, I want to thank Frank for always being accessible and ready to share your thoughts. I would also like to thank everyone at MMIV for making me feel welcome, and for sharing helpful ideas throughout the process.

I would also like to express my gratitude to my supervisors, for granting me the opportunity to attend the annual Nordic ISMRM conference in Lund (travel costs covered by a research grant from the Norwegian Cancer Society to Renate Grüner), where I got to present a poster during the initial phase of my thesis writing. Additionally, I had the privilege of presenting another poster at the MMIV conference in December. These experiences provided me with good feedback from some of the most eminent researchers in the field.

Thanks to Fulvio Zaccagna and Ferdia A. Gallagher (University of Cambridge, UK) for sharing the MRI data for this project, and thanks to Andrew Priest (Addenbrooke's University Hospital, Cambridge, UK) for helping with the data collection.

Furthermore, I would like to express my appreciation to all my fellow students throughout the four years I have had at the university, for always being supportive of each other, and for sharing helpful ideas in the process of writing the thesis. Last but not least, I would like to thank my family and my partner for all the support, and for giving me feedback on my work.

Thea Våge

Bergen, June 2023

Nomenclature

Abbreviations

ANN Artificial Neural Network

CHARMED Composite Hindered and Restricted Model of Diffusion

CPU Central Processing Unit

DL Deep Learning

DMIPY Diffusion Microstructure Imaging in Python

dMRI Diffusion Magnetic Resonance Imaging

EES Extracellular-Extravascular Space

ELU Exponential Linear Unit

EPI Echo Planar Imaging

FA Fractional Anisotropy

FFNN Feed Forward Neural Network

FID Free Induction Decay

GPU Graphics Processing Unit

GRE Gradient Echo Sequence

GT Ground Truth

H&E Hematoxylin and Eosin

HGG High Grade Glioma

IC Intracellular

IVIM Intravoxel Incoherent Voxel Model

LGG Low Grade Glioma

MDT Maastricht Diffusion Toolbox

MIX Microstructure In Crossings

ML Machine Learning

MLP Multilayer Perceptron

MR Magnetic Resonance

MRI Magnetic Resonance Imaging

MSE Mean Squared error

NMR Nuclear Magnetic Resonance

NN Neural Network

PGSE Pulsed Gradient Echo Sequence

RELU Rectified Linear Unit

RF Radio Frequency

ROI Region of Interest

SE Spin Echo

SGD Stochastic Gradient Descent

SNR Signal to Noise Ratio

TE Echo Time

VASC Vascular

VERDICT Vascular, Extracellular and Restricted Diffusion for Cytometry in Tumors

Symbols

α Flip angle

β Exponential decay rate

Δ	Diffusion duration
δ	Diffusion time
η	Viscosity
γ	Gyromagnetic ratio
$\hat{y}$	Predicted value
$\hbar$	Plancks constant
λ_n	Eigenvalues
λ_{iso}	Extracellular-extravascular diffusion coefficient
μ	Magnetic moment
$\overrightarrow{B}$	Magnetic field strength
$\overrightarrow{M}$	Net magnetization
θ,ϕ	Orientation angles
b	b-value: degree of diffusion weighting
B_0	Main magnetic field
B_1	Oscillating magnetic field
D	Diffusion coefficient
f_{EES}	Extracellular-extravascular volume fraction
f_{IC}	Intracellular volume fraction
f_{VASC}	Vascular volume fraction
G	Gradient strength
k	Boltzmanns constant
S	MR signal
S_0	MR signal at baseline
T	Tesla

t	Time
y	True value
R	Radius

List of Figures

List of Tables

Contents

Chapter 1

Introduction

Gliomas belong to a group of tumors in the central nervous system and are the most prevalent primary brain tumor among adults [1] [2]. These tumors are characterized by poor prognosis as the glioma often displays a significant histopathological heterogeneity, which provides poor therapy response and frequent tumor recurrences [3] [4]. The choice of treatment is today decided through a biopsy, which is invasive and does not capture the full heterogeneity in the tumor [5].

The lack of effective treatments is a major factor contributing to the low prognosis, and the progress in treatment is also limited by difficulties with monitoring clinical treatment response. As of now, histopathology is employed to monitor biomarkers and it is often considered the reference standard of treatment response. This approach has several challenges as there is a paucity of evidence that earlier diagnosis influences the prognosis. Additionally, the invasive biopsy technique with repeated tissue sampling has a high risk of morbidity in patients with glioma compared to other systemic cancers [6]. Repeated biopsies for treatment monitoring are not so often done, but another used technique for treatment monitoring is structural magnetic resonance imaging (MRI).

Since structural MRI scans are already time- and resource-consuming, it could be beneficial to include other analyses within the same scan. There exist different models that infer tissue microstructure from diffusion MRI (dMRI) scans [7] [8] [9] [10] [11] [12]. Analysis using such models can give measures of tissue parameters such as tumor cell size and volume fractions of different compartments within the glioma tissue. As an MRI scan is required nevertheless, data for such dMRI analysis could easily be added to the imaging protocol. However, a requirement for this to be clinically applicable would be that it is fast in both acquisition and analysis, in addition to providing relevant additional information.

1.1 Recent Developments and Contribution

In 2014, Panagiotaki et al. presented a new technique called vascular, extracellular and restricted diffusion for cytometry in tumors (VERDICT), to quantify and map histological features of tumors from dMRI [12]. The VERDICT model was successfully used to differentiate tumors from benign areas in prostate tissue and to reveal differences in microstructural parameters such as volume fractions, indicating that the model had clinical potential [13]. Following this work, Zaccagna et al. demonstrated that VERDICT parameters of tissue structure in gliomas correlated with parameters derived from histology. They stated that this method shows promising results for diagnostic stratification and treatment response monitoring of gliomas in the future.

The conventional approach for applying the VERDICT model to data has often involved utilizing either an iterative optimization process with a non-linear least squares (NLLS) optimization [14] [15], or employing a convex optimization method [16]. The NLLS optimization is slow and its performance is easily affected by different local minima in the fitting objective function [17]. For this reason, deep learning techniques have recently been proposed to overcome this problem. Self-supervised deep learning networks have been applied in the fitting of several different microstructure models. However, it is yet to be applied for VERDICT fitting of glioma tissue.

1.2 Problem Statement and Thesis Outline

Effective techniques such as the VERDICT method for non-invasive measurement of tumor cellularity exist, but they are hampered by both long acquisition and analysis time. Therefore, the purpose of this thesis was to explore whether the VERDICT microstructure fitting in glioma tissue could be approached with the novel self-supervised fully connected neural network technique. The aim was to investigate the potential of the deep learning approach both in terms of computational time and accuracy compared to existing methods. This was done in a systematic manner, by first constructing a neural network and assessing it on a large-scale dataset. Secondly, clinical MRI data was mimicked and the approach was tested on *in vivo* data. As a third and explorative step, the effect of data collection schemes was investigated. Finally, the model was compared to existing methods in terms of computational time and accuracy. Results from the investigation of these problems can facilitate realistic integration into clinical imaging, and thus facilitate the diagnostics and treatment monitoring of gliomas.

Chapter 2

Theory

2.1 Magnetic Resonance Imaging

Magnetic resonance imaging (MRI) is a powerful modality used for medical diagnostics. The use of magnetic fields enables the generation of anatomical images of human tissue, as well as functional images illustrating blood flow or diffusion. In the following section, the basic principles of MRI and dMRI are presented, and details can be found in Cercignani et al. [18], Brown et al. [19] and Grüner [20].

2.1.1 Basic Principles of MRI

Spin and Magnetic Momentum

The basic idea of MRI is to image the proton interactions in the body, by exploiting a basic property of the protons called nuclear spin. This is a fundamental propriety such as a particle's mass. The spin properties of an atom are dependent on the composition of nucleons in an atom. The spin quantum number associated with protons is s = 1/2. Associated with the spin is a magnetic moment μ, which is proportional to the atom's spin:

$$\mu = \gamma S \tag{2.1}$$

In this equation, the gyromagnetic ratio, γ, relates the momentum to the atoms spin. The hydrogen atom has the highest gyromagnetic ratios, gamma = 42.57 MHz/T. This is an advantage in addition to that the human body contains a lot of hydrogen, as they also provide the strongest magnetic proprieties and therefore the highest nuclear magnetic resonance (NMR) signal. The magnitude of this magnetic moment is well-known:

$$|\mu| = \gamma\hbar\sqrt{I(I+1)} \tag{2.2}$$

, where I is the spin number = 1/2 for protons, γ is the gyromagnetic ratio and $\hbar$ is Plancks constant. When these nuclei are placed in a magnetic field, they will align parallel or antiparallel with the field, where the magnetic momentum in the direction of the field is described as:

$$\mu_z = \pm\frac{1}{2}\gamma\hbar \tag{2.3}$$

, where the positive sign represents the momentum in the z-direction while aligned with the magnetic field, and the negative sign represents when aligned against the field. These two orientation possibilities result in that one can have two different energy states. The magnitude of this individual magnetic momentum is very small, and thus is it the net magnetization that is measured in MRI. This can be explained as an ensemble of N spins with magnetic momentum μ:

$$\overrightarrow{M} = \sum_{n=1}^{N} \mu_n \tag{2.4}$$

The behavior of this net magnetization is described by the fundamental physics of the Bloch equation, which lays the fundamental of the measurement of MRI signal as the solution of the Bloch equation is the signal equation. The equation of motion is simplified as the following equation as a function of the magnetic field $\overrightarrow{B}$:

$$\frac{d\overrightarrow{M}}{dt} = \gamma\overrightarrow{M} \times \overrightarrow{B} \tag{2.5}$$

Excitation and Relaxation

In order to measure the MRI signal, the net magnetization needs to experience a change over time. By adding an oscillating field B_1, the magnetization is perturbed away from its positioning along B_0. This pulse, which is often called a radio frequency (RF) pulse, is applied perpendicular to B_0 and is typically applied for a few micro and milliseconds. Immediately after the application of the RF pulse, the spin system sends out a measurable free induction decay signal (FID). In order to spatially encode this signal it is delayed and the signal is measured as an echo. The time between the FID and the readout of the signal is called echo time (TE).

Immediately after the excitation, the net magnetization returns towards its original

position along the direction of the B_0 field, now called the z-axis. The relaxation is explained in the less simplified version of the earlier introduced Bloch equation introducing the relaxation terms T_1 and T_2:

$$\frac{d\vec{M}}{dt} = \gamma\vec{M} \times \vec{B} - \frac{M_x\,\vec{i} - M_y\,\vec{i}}{T_2} - \frac{(M_z - M_z^0)\,\vec{k}}{T_1} \tag{2.6}$$

M_z^0 is the net magnetization along the z-axis when only the B_0 field is present, and is more precisely called the thermal equilibrium magnetization. T_1 is a time constant describing the longitudinal relaxation process along the z-axis, more precisely it measures how fast 63% of the net magnetization is regained. The T_2 time constant describes the transversal relaxation process in the xy-plane and gives a measure of the time when 37% of the transversal magnetization is left.

Image Formation

These time constants, T_1 and T_2 are typically referred to as the contrast of an MRI, along with other used contrasts such as diffusion weighting that is central in this project. This means that the image is weighted to show the longitudinal or transversal tissue properties. In order to construct an image with contrast, the pulse sequence parameters are manipulated. These parameters include α, the flip angle of which the RF pulse perturbs the magnetization vector, TR which is the time delay between to successive RF pulses, TI which is the time between pulses to selectively zero out the signal contribution from a particular tissue and the TE.

Along with these pulse sequence parameters, gradient magnets are used to create smaller fields that can give variations in all directions. This is the central part of how a specific image slice is excited, as a gradient is applied in the z-direction along the B_0 field creating variation in the Larmor frequencies. As this gradient is applied during an RF-pulse, only a bandwidth matches the Larmor frequency and it is excited. During readout, gradients in x- and y-direction are applied to allow frequency-encoding and phase-encoding, respectively. As the signal is acquired, the essential concept *k-space* is introduced to construct an image. This can be seen as a data matrix where the measured spatial frequencies are stored before it is converted to spacial coordinates with a Fourier transform.

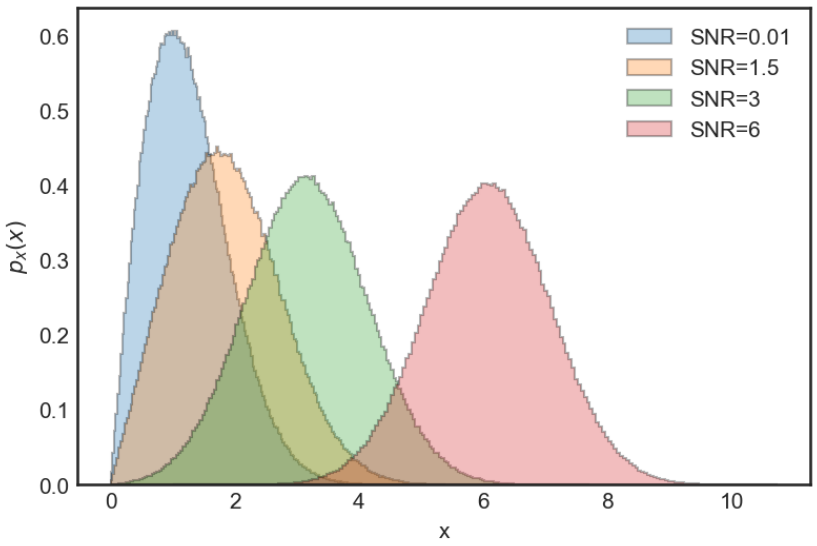

Figure 2.1: Rice distribution at varying signal-to-noise ratios (SNR).

Noise

In the constructed images, measurement of signal intensity at a low signal level is compounded with its noise [21]. Such noise is shown to be governed by a Rician distribution [22]. In the presence of Rician noise, the probability distribution for a measured pixel intensity I_0 is given by 2.7 [23]. From plot 2.1 it can be seen that the Rice distribution is far from being Gaussian for a small signal-to-noise ratio (SNR), but as the ratio increases, the distribution approximates the Gaussian distribution.

$$f(x|\nu,\sigma) = \frac{x}{\sigma^2} e^{-(x^2+\nu^2)/2\sigma^2} I_0\left(\frac{x\nu}{\sigma^2}\right) \tag{2.7}$$

Sequences

Several methods exist to sample the MRI signal using so-called pulse sequences, which are designed with a certain order and intervals of the RF pulses. One of the two most used sequences is the gradient echo (GRE) sequence, which samples one line of k-space for each RF pulse. The x and y-gradients are varied both before and during a readout to move in the k-space. Secondly, the spin echo (SE) sequence is often used. This sequence includes a second RF pulse of 180 degrees if the α is 90 degrees. Other sequences such as echo-planar imaging (EPI) can sample the whole k-space in one RF pulse. This method is often applied for diffusion-weighted imaging.

2.1.2 Diffusion MRI

Diffusion

In 1905, Albert Einstein established the existence of molecules while studying Brownian motion, when he observed that the fundamental particle property was not its velocity but its displacement [24]. He then showed that the mean displacement in one direction of a particle undergoing so-called free diffusion is given by expression 2.8, where t is time and the diffusivity in m^2/s. D is given by expression 2.9, where k_b is the Boltzmann constant, t is time, η is the viscosity and R is the particle radius.

$$< r^2 >= 6Dt \tag{2.8}$$

$$D = \frac{k_b t}{6\pi\eta R} \tag{2.9}$$

The easiest way to describe free diffusion is as the random motion of molecules, which means that the diffusion is independent of direction [25]. Such diffusion is often referred to as isotropic diffusion. However, for imaging of diffusion, we consider it more interesting to study the presence of boundaries in water diffusion. When water is no longer diffusing with random motion, it allows us to measure the geometrical properties of the boundaries in tissue [25].

In order to include information on the surrounding tissue structures, two new terms are introduced: restricted and hindered diffusion. Restricted diffusion is what occurs when water molecules are trapped inside a cell, which is why this motion is also referred to as intracellular diffusion [25]. On the other hand, hindered diffusion is when the moving water molecules are impacted by the geometrical structures, as they meet obstacles. This is also referred to as extracellular diffusion as it is present in the space between cells [25]. The assumed behavior of cells within a restricted and a hindered compartment is illustrated in Fig 2.2.

In addition to these intracellular and extracellular compartments, other properties of diffusion are of interest in imaging of diffusion. When the water molecules are moving in a cylindrical boundary as they often do in tissue, the water diffusion becomes

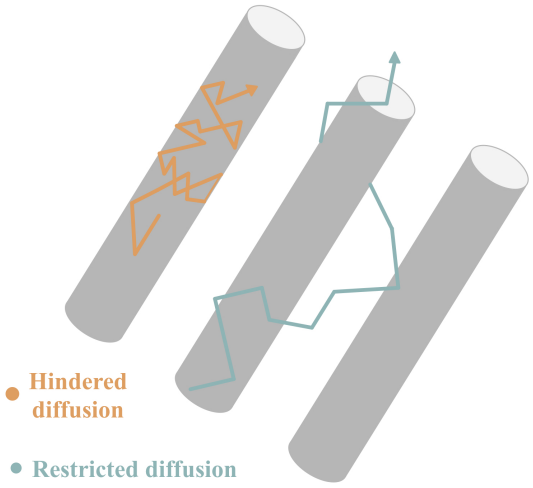

Figure 2.2: Illustration of the hindered and restricted diffusion process.

dependent on the cylindrical boundaries [25]. When a water molecule is directionally dependent on the cylindrical shape, it is referred to as "anisotropic", while the opposite is called "isotropic" diffusion [25].

Imaging of diffusion

Since diffusion is defined as random motion on molecular levels, one image voxel will contain many diffusion paths. This means that an MRI scan will contain information about diffusion contributions, which could reveal tissue information. In 1965, Stejskal and Tanner introduced the pulsed gradient spin echo (PGSE) sequence [26], which establishes the fundamentals of modern diffusion imaging. This sequence as illustrated in figure 2.3 applies symmetric and strong gradients on each side of the 180°-pulse. The concept of these pulsed gradients is that the phase shift induced by the first pulse is then reversed by the second pulse for all spins that are stationary. However, for spins that have diffused and changed location in between the two pulses, they will be out of phase and contribute with less signal to be measured.

This image formation makes the measure sensitive to several acquisition parameters: the diffusion time between the two gradients, Δ, the duration of the gradient, δ, and the strength of the Gradient, G. Now, taking this into account, a diffusion term was added to the previously defined Bloch equation [27], as demonstrated in equation 2.10.

$$\frac{d\vec{M}}{dt} = \gamma\vec{M} \times \vec{B} - \frac{M_x\,\vec{i} - M_y\,\vec{i}}{T_2} - \frac{(M_z - M_z^0)\,\vec{k}}{T_1} + \nabla * D\nabla M \tag{2.10}$$

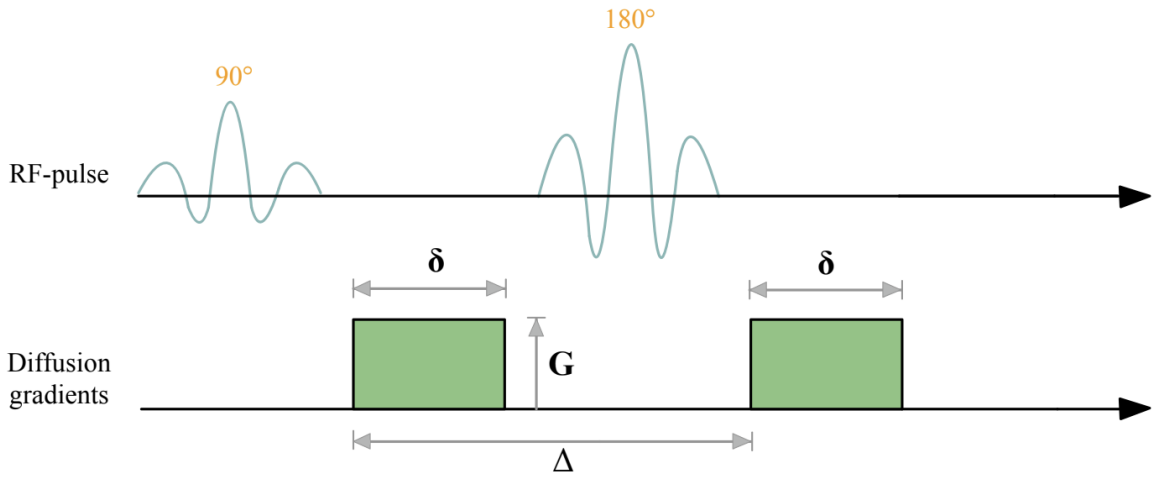

Figure 2.3: A pulsed gradient spin echo (PGSE) sequence with an RF-pulse and a diffusion sensitizing gradient pair, with gradient strength G, gradient duration δ and diffusion time Δ.

In practice, this inclusion of the diffusion terms gives an echo attenuation according to Stejskal and Tanner [26] as described in expression 2.11. This equation is only derived for isotropic diffusion, so no three-dimensional motion is measured, as the diffusion coefficient D is a scalar only.

$$A(TE) = e^{-\gamma^2 G^2 \delta^2 D(\Delta - \delta/3)} \tag{2.11}$$

The gradient pulses are applied to obtain diffusion-weighted images, but in order to quantify the findings it is necessary to further determine the degree of diffusion weighting. As the solution of the extended Bloch equation 2.10 developed by Torrey [27] becomes complex, it was suggested to summarize all the gradients effects by Le Bihan and Breton in 1965 [28]. They introduced the term that is generally known as the b-value as expressed in 2.12. With this factor, the signal attenuation expression can be rewritten to a simpler expression as can be seen in 2.13.

$$b = \gamma^2 \int_0^{TE} [\int_0^t \mathbf{G}(t')dt']^2 dt \tag{2.12}$$

$$A = e^{-bD}, \qquad b = -\gamma^2 G^2 \delta^2 (\Delta - \delta/3) \tag{2.13}$$

This b-value is now a measure of how much diffusion weighting there is in a sequence, based on the acquisition parameters γ, G, δ and Δ. With this attenuation, the diffusion signal obtained for such a PGSE sequence dependent on the b-value and diffusion

coefficient is expressed in equation 2.14.

$$S(b) = S_0 e^{-bD} \tag{2.14}$$

The b-value is in a standard SE-sequence with no additional pulsed gradients around $1 s/mm^2$ [7], which means that the diffusion effects are negligible. However, with the use of gradient pulses, these effects are no longer negligible and may lead to the misdetermination of the b-value. This means that more quality-assuring steps are necessary for the imaging of diffusion. One problem is that when these diffusion gradients are applied, cross-terms can be produced by the combination of the imaging and diffusion gradient pulses [24], which might lead to diffusion-related attenuation effects [29]. Especially can Eddy currents extend over the imaging readout gradient pulse, which has the potential to vary the b-value and therefore lead to incorrect estimation of the diffusion coefficient.

High b-values and noise

A high b-value can in some cases be favorable. However, high b-values come at the cost of lower SNR. A study at 3T has found that SNR at b = 3000 s/mm^2 was 30.8 $\pm$21.7, which was approximately one-third of the SNR at b = 1000 s/mm^2, 93.5 $\pm$42.1 [30].

2.2 Diffusion microstructure imaging

The standard diffusion-weighted imaging techniques can provide useful information about tissue structure, but it has in some studies shown that there is no correlation between the signal value and tumor cellularity [31] [32]. Thus, new modeling approaches have been developed in order to further understand the composition of the tissue. In microstructure imaging, the signal is modeled as arising from different compartments. Such model-based DWI can provide more information that the signal acquired from the imaging sequence, and that can provide many advantages within glioma treatment response assessment.

2.2.1 Diffusion Tensor Imaging

Traditionally, diffusion microstructure imaging has been approached with the simple diffusion tensor model for relating the diffusion MRI to the tissue diffusivity [33]. This method called diffusion tensor imaging (DTI) is based on the assumption that local diffusion can be described with a three-dimensional Gaussian distribution, and that this distribution has a covariance matrix that is proportional to the diffusion tensor, **D**. This diffusion tensor adds a new dimension of complexity from the original signal equation from a diffusion sequence as described in expression 2.14, which makes it accountable for molecular mobility, and gives information about direction and correlation between directions [34]. Seen from a reference frame $[x', y', z']$ according to the diffusion, the diffusion signal is now represented as in the following equation, rather than the simple form explained in formula 2.14 in section 2.2.

$$S(b) = S_0 e^{-b_{xx'}D_{xx'} - b_{yy'}D_{yy'} - b_{zz'}D_{zz'}} \tag{2.15}$$

This way of representing the diffusion coefficient provides more information, as eigenvectors and eigenvalues can be used to estimate a fractional anisotropy (FA), which is a measure of how asymmetric the diffusion within a voxel is. For a diffusion matrix **D** with eigenvalues λ_1, λ_2 and λ_3, the fractional anisotropy is measured according to equation 2.16.

$$FA = \sqrt{\frac{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_1 - \lambda_3)^2}{2(\lambda_1^2) + \lambda_2^2 + \lambda_3^2}} \tag{2.16}$$

A map of the derived FA-values can be constructed to represent a diffusion tensor image. This resulting image has a gray-scale map where brighter areas contain tissue with higher anisotropy.

2.2.2 Biophysical Models and Multi-Compartment Models

This traditional DTI method can provide some useful biomarkers, but is in most cases not enough to describe the composition of the tissue. As the diffusion signal is modeled as arising from different compartments, just one single compartment might not be enough in all cases. The DTI model assumes free diffusion in all of the tissue and does not take into account restricted and hindered diffusion. Newer methods have shown the possibility of fitting more complex models that provides more detailed modeling of diffusion in tissue. These methods can provide information such as cell size and density directly from the diffusion MRI signal. Such models are often so-called multi-compartment models, such as the one modeled by Le Bihan et al. in 1988 [7]. This research introduced the representation of tissue as separate diffusive processes with a two-compartment model that distinguishes diffusion and perfusion in tissue. This intravoxel incoherent voxel model (IVIM) models the signal by representing the tissue as consisting of two compartments: one vascular to represent the pseudo-diffusing water within blood vessels, and one non-vascular compartment which represents the diffusion of water in and around cells.

Following the work of Le Bihan et al. [7], several different models have been developed. One of the simplest methods is the well-known ball-stick compartment model which combines a ball and a stick compartment to separate intra-axonal from extra-axonal water signals [9]. A new model combining hindered and restricted models of water diffusion (CHARMED) combines a so-called hindered extra-axonal compartment described by a diffusion tensor with a restricted intra-axonal compartment within cylinders. The use of cylinders as a modeling technique allows for measurements of cell radius.

Comparison studies of compartment models in the brain have shown that three-compartment models best explain the data [14] [35]. Furthermore, especially models combining diffusion tensor models with extracellular space and a cylindrical intracellular component with a spherical compartment with a non-zero radius.

2.2.3 VERDICT Model

A multi-compartment model for use in tumorous tissue has been proposed to overcome the complex challenges as described. The model called vascular, extracellular and restricted diffusion for cytometry in tumors (VERDICT) models the water diffusion in tissue by the design including three different compartments [12]. As previously demonstrated as a good compartment modeling combination [14] [35], this model combines two diffusion tensor models with a spherical compartment. Tumor cells, the hindered extracellular and the vascular space are modeled as a restricted sphere, an isotropic Gaussian ball and a stick compartment, respectively. The VERDICT signal is the sum of these three models, which each describe the diffusion MR signal in its respective cellular environment, 2.17.

$$E_{\text{VERDICT}} = \underbrace{f_{\text{IC}}\overbrace{E(D|\lambda_{\text{intra}})}^{\text{Sphere}}}_{\text{S1}} + \underbrace{f_{\text{EES}}\overbrace{E(\cdot|\lambda_{iso})}^{\text{Ball}}}_{\text{S2}} + \underbrace{f_{VASC}\overbrace{E(\lambda_{\parallel},\mu)}^{\text{Stick}}}_{\text{S3}} \tag{2.17}$$

These three model compartments are derived from the tissue and give three different signals. The S_1 signal from the sphere compartment comes from intracellular (IC) water trapped inside cells. The S_2 signal comes from water in the extracellular-extravascular space (EES) adjacent to, but outside the tumor cells and blood vessels, while the S_3 signal comes from water in the blood in the capillaries. The VERDICT model requires a signal sampled with δ unequal to Δ and G as opposed to simpler methods such as DWI, where only G is varied to change the b-value.

Intracellular Compartment

The S_1 signal comes from the IC compartment, which is demonstrated as a spherical compartment in the original VERDICT paper [12]. The IC signal can be modeled using various models of restricted diffusion in domains with different shapes such as cylinders or ellipsoids, but spheres are good representations in the case of isotropic tumor cells. Balinov et al. [36] presented a model for the signal from the spherical compartment 2.18, where $j_n(x)$ is the spherical Bessel function of the first kind and α_{nm} is the mth

root of the equation $j'_n(\alpha) = 0$.

$$E(\delta,\Delta,g) = \frac{9[\gamma G\delta R cos(\gamma G\delta R) - sin(\gamma G\delta R)]^2}{(\gamma G\delta R)^6} - 6(\gamma G\delta R)^2 \sum_{n=0}^{\infty} [j'_n(\gamma G\delta R)]^2 \sum_m \frac{(2n+1)\alpha_{nm}^2}{\alpha_{nm}^2 - n^2 - n} \times exp(-\frac{\alpha_{nm}^2 D\Delta}{R^2}) \frac{1}{[\alpha_{nm}^2 - (\gamma G\delta R)^2]^2} \tag{2.18}$$

However, this approximation is only exact for very short gradient pulses. As this is not always applicable, Neuman et al. [37] discuss that phases are Gaussian distributed. Following this work, Murday and Cotts derived a signal equation from the investigation of the accuracy of the Gaussian phase-distribution (GPD) approximation for spins confined to spheres [38], where the expression of the echo attenuation for the GPD approximation in the case of the sphere is given as represented in equation 2.19.

$$ln[E(\delta,\Delta,g)] = -\frac{2\gamma^2 g^2}{D} \sum_{m=1}^{\infty} \frac{\alpha_m^{-4}}{\alpha_m^2 R^2 - 2} \times [2\delta] - \frac{2 + e^{-\alpha_m^2 D(\Delta-\delta)} - 2e^{-\alpha_m^2 D\delta} - 2e^{-\alpha_m^2 D\Delta} + e^{-\alpha_m^2 D(\Delta+\delta)}}{\alpha_m^2 D} \tag{2.19}$$

, where α_m is the mth root of the equation $[1/(\alpha R)]J_{3/2}(\alpha R) = J_{5/3}(\alpha R)$, and D is the IC diffusion coefficient and R is the cell radius.

The DMIPY modeling framework [39], which will be further explained in section 2.2.4, have implemented this GPD to model the spherical compartment. The modeling of the spherical compartment is based on the developed equation from Murday and Cotts in 1968, presented in equation 2.19. For simplification, their implementation contains the 100 fist roots of α_m as constants. The results of the DMIPY model fitting reflects what is originally found in the VERDICT paper [12].

Extracellular-Extravascular Compartment

The S_2 signal, which comes from the EES compartment uses a diffusion tensor model as previously explained. For an isotropic arrangement of cells, this can be constrained to be isotropic, and anisotropic in the case of anisotropic cell arrangement. In the VERDICT demonstration [12], the isotropic arrangement is assumed and the compartment is modeled as a geometrical ball. The DMIPY framework uses this geometrical ball with

Gaussian distribution to model the compartment. The signal is calculated as shown in expression 2.20.

$$E(\lambda_{iso}, b) = e^{\lambda_{iso} * b} \tag{2.20}$$

, where λ_{iso} is the EES diffusion constant, and b is a vector consisting of the acquisitions b-values.

Vascular Compartment

The S_3 signal from the blood in the capillaries is referred to as the vascular compartment. This compartment also uses a diffusion tensor model to represent the vascular pseudo-diffusion. If the data contains a high degree of anisotropy in this component, the pseudo-diffusion can be oriented along a single direction, but more isotropic models might be more suitable in situations where this is not the case. In the DMIPY framework, they use this diffusion tensor model that assumes a zero radius. This resulted in the stick model [9], whose signal can be modeled with expression 2.21.

$$E(b, \lambda_{||}, \theta, \phi) = e^{b * \lambda_{||} (\theta \cdot \phi)^2} \tag{2.21}$$

, where $\lambda_{||}$ is the vascular pseudo-diffusion constant, and θ and ϕ are the orientation angles.

The VERDICT Modeling

A total of seven parameters is included in these three compartments: the three-volume fractions from each compartment f_{IC}, f_{EES}, f_{VASC}, the IC diffusion constant, d_{ic}, the EES diffusion constant, d_{ees}, the pseudo-diffusion coefficient $\lambda_{||}$, and the spherical radius, R. The VERDICT model assumes that the vascular compartment consists of cylinders with uniformly distributed orientations, and thus the two orientation parameters, θ and ϕ are constant. Figure 2.4 from Panagiotaki et al. [13] illustrates these three compartments in prostate tissue.

2.2.4 Anatomy of Existing Multi-Compartment Modeling

There are a few existing frameworks that are publicly available that support the microstructure fitting of the VERDICT model. One of these frameworks is the DMIPY model.

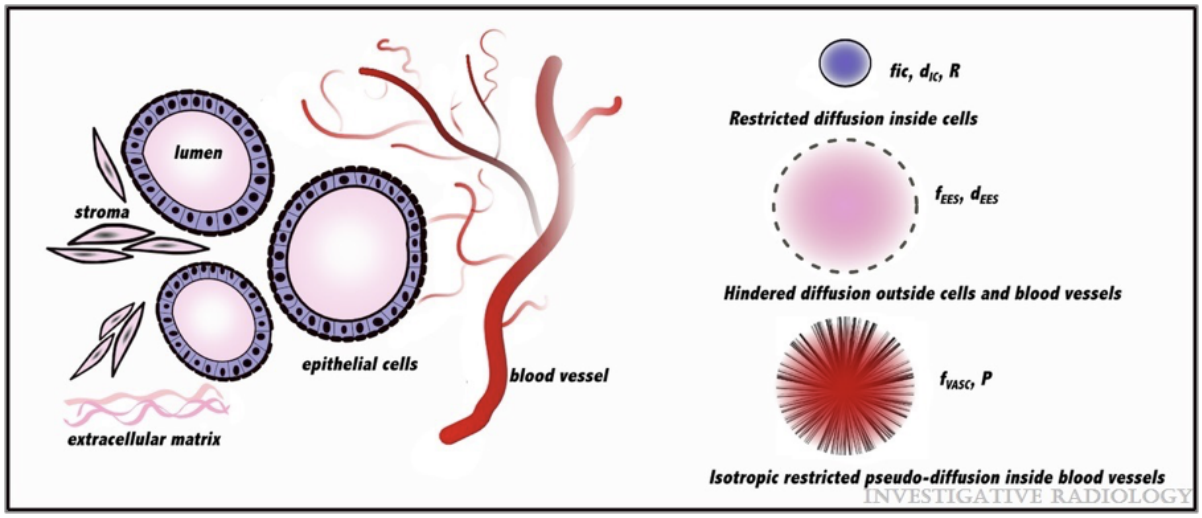

Figure 2.4: Figure from Panagiotaki et al. [13]. Schematic representation of the VERDICT model in tissue. Intracellular (IC) diffusion, d_{ic}, IC volume fraction, f_{IC} and radius R are represented in the restricted diffusion inside cells. Extracellular-extravascular (EES) diffusion constant d_{EES} and volume fraction f_{EES} are represented in the hindered diffusion inside cells and blood vessels. Pseudo diffusion inside blood vessels, P (in this project also called $\lambda_{||}$), and vascular volume fraction f_{VASC} are represented in the isotropic restricted pseudo-diffusion inside blood vessels.

Diffusion Microstructure Imaging in Python Toolbox

The Diffusion Microstructure imaging in Python (DMIPY) toolbox was presented to provide deeper insights into the capabilities of MC-modeling [39]. How the DMIPY framework deals with the three compartments of the VERDICT model has already been explained, but how the model optimizing works is actually done remains to be explained. Once the MC-model has been defined as explained in section 2.2.3, the framework provides a way to link, fix or optimize parameters. Especially relevant for the fitting of the VERDICT is the fixing of parameters, which implies fixing a known global that has previously been optimized. Once this is done, DMIPY provides two different optimization strategies for fitting of MC-models. One of the two is the microstructure in crossings (MIX) algorithm [40].

This MIX algorithm is designed to return accurate results when applied to highly complicated MC models such as VERDICT. By exploiting a separable structure in the optimization problem, this method allows estimating detailed axonal features in the presence of multiple fiber orientations [40]. However, its accuracy comes at the cost of being slow and is not preferred for simpler models [39].

Maastricht Diffusion Toolbox

The Maastricht diffusion toolbox (MDT) is a toolbox that allows for VERDICT modeling [15]. This toolbox uses the Powell optimization routine as standard, which is a derivative-free nonlinear local optimization algorithm [41]. This optimization routine was shown to outperform the other algorithms in terms of running time, fit accuracy and precision. The algorithm commences the search by initializing n search vectors, each of which comprises all n model parameters. At each iteration, the algorithm updates each search vector individually by carrying out a line search along each dimension. In this manner, the algorithm generates vector combinations and substitutes outdated combinations until a stopping criterion is fulfilled. The model and optimization algorithm is automatically compiled into OpenCL code that can be executed on both a central processing unit (CPU) and graphical processing unit (GPU) [42].

2.2.5 Diffusion Microstructure Imaging in Gliomas

This three-compartment biophysical model, VERDICT, has previously been used to provide clinically useful cancer biomarkers from various tissue structures [13] [43] [44]. These applications demonstrate that the three compartments previously described in the VERDICT model are aiming to represent the tumor tissue structure. The IC compartment with restricted diffusion is expected to be the greatest in the tumor cells. Furthermore, the radius is expected to be higher in a higher-grade glioma (HGG) compared to a lower-grade glioma (LGG). Fig. 2.5 from Zaccagna et al. [43] illustrates the cell radius differences in LGG and HGG and the hematoxylin and eosin (H&E) stain from LGG and HGG tissue.

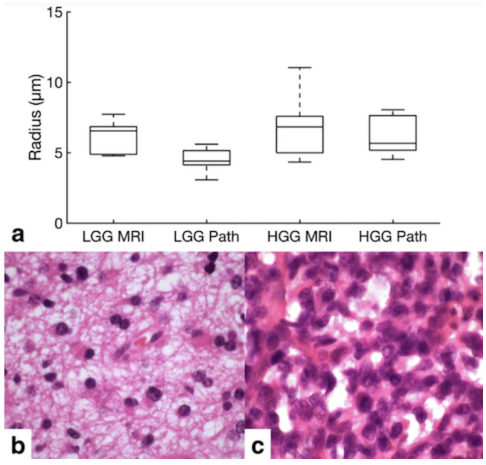

Figure 2.5: Illustration from Zaccagna et al. [43]. (a) shows box and whisker plots of the cell radius measured by VERDICT MRI and the pathology examination for LGG and HGG. (b) and (c) showing H&E stains from LGG and HGG, respectively.

Prior studies on glioma-specific cases have given estimations of the VERDICT parameters [44] [43]. Roberts et al. estimate the cell radius of gliomas to be $10.6 \pm 0.6\ \mu m$, while Zaccagna et al. estimate it to be $6.8 \pm 2.3\ \mu m$ for HGG and $6.7 \pm 1.2\ \mu m$ for LGG. The vascular volume fraction was shown to be significantly lower than the IC and EES fractions. Roberts et al. estimate the IC, EES and vascular volume fractions to be approximately 0.2, 0.75 and 0.05, respectively, while Zaccagna et al. estimate the IC and EES fractions to be 0.13 and 0.88 for HGG, and 0.08 and 0.92 for LGG, respectively. As shown in figure 2.5 (a), Zaccagna et al. also did a comparison to histological measures that verifies that the estimated values are within expected ranges.

2.3 Machine Learning

Tom Mitchell gave a description of the task machine learning (ML) is trying to solve as follows: A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance a task in T, as measured by P, improves with experience E [45]. In recent years, a renewed interest in machine learning is seen, due to new algorithms, more computational performance and data. Machine learning now has a central role in tasks like speech recognition, translations and product recommendation. This advancement has led to a rise in the use of AI in radiological imaging tasks such as reconstruction, segmentation, diagnosis, prognosis and risk assessment [46].

2.3.1 Machine Learning Fundamentals

Typically, an ML process is an approach in three steps that distinguishes the phases of training, testing and validating. The training phase is when the ML algorithm is given data, and the set of weights or decision points for the model is updated until it reaches its optimal performance [47]. ML algorithms usually address two main types of problems: classification problems and regression problems. In classification problems, the model analyzes features in the input data and predicts the category to which the input data belongs. In regression problems, labels are represented by arbitrary numerical values instead of categories. The objective is to develop a model that can make predictions that closely resemble the actual label values [48].

Machine learning adopts three conventional approaches for the training process: supervised, unsupervised and reinforcement learning. Supervised learning is an approach where the model is given the ground truth, often called labels. The goal of the model is then to make as good predictions as possible to match the given ground truth values. When the model is not given the ground truth during training, this is called unsupervised learning. In this approach, the algorithm is looking for patterns and relations from the unlabelled data. A third approach is reinforcement learning, which is a combination of the two above. This method begins with a classifier using labeled data, before it tries to improve the predictions [47].

Self-Supervised Learning

The branch called unsupervised learning contains many different approaches to solve the handling of unlabeled data. In this project, a self-supervised approach is being explored. Self-supervised learning is known as a general method that uses the data itself and infers supervisory constraints on the data to learn image features [49]. Previous studies have shown the potential of using MRI signal equations as such supervisory constraints, where the model is learning by simulating signals from the learned features in the data [49] [10] [50] [51] [52].

2.3.2 Deep Learning

In some cases where traditional machine learning comes to short, deep learning (DL) has been making major advances in solving problems that have not yet been successfully solved. DL has been shown to be good at discovering intricate structures in high-dimensional data [53]. This has resulted in its good accuracy in tasks such as image recognition[54] [55] [56] [57], reconstructing brain circuits [58] and natural language processing tasks.

DL is a sub-field of ML, where the representation of data is turned into several representation layers [59]. These representations of data are in DL learned through models called neural networks (NN), which consist of several stacked layers. Each layer deals with the data by creating weights which are numbers that specify what to do with the input data, where the learning process is based upon finding the weights that provide the best fit to the target.

Training of Neural Networks

For the network to control the weights and make sure they learn in the direction of the targets, three fundamental terms are introduced: the loss function, optimizer and backpropagation. The loss function computes the score of how well the predictions do compared to the true target, and this score is used as feedback in order to adjust the weights in the right direction. The optimizer implements this loss function and does what is called the backpropagation algorithm. So once the data has been fed forward through the neural network and the loss has been calculated, the loss is then propagated backward through the top layers and the chain rule is applied. The chain rule is derived from calculus and gives us the relation shown in formula, which tells us

how to differentiate composite functions such as a function $h(x) = f(g(x))$ 2.22 [60]. In this context, the chain rule computes the amount of contribution each parameter had to the loss value [59].

$$\frac{d}{dx}[f(g(x))] = f'(g(x)) * g'(x) \tag{2.22}$$

The optimizer uses a gradient vector to properly adjust the weight vector. This gradient vector is an indication of how much the error would increase or decrease if a specific weight was increased by a small amount before the weight vector is then adjusted in the opposite direction [53]. In most cases, this procedure is done by implementing a variant of stochastic gradient descent (SGD). This method is a drastic simplification of the traditional gradient descent method because it picks a random example of data points from the training batch and computes the gradient based on that small set [61].

One of the common variants of SGD is the Adam optimizer [62]. This method is combining the advantages of earlier, well-known methods: AdaGrad [63] which has the advantage of working well with sparse gradients, and RMSProps which works well in on-line and non-stationary settings [62]. This method computes individual adaptive learning rates for different parameters from estimates of the first and second moments of the gradients. Given a step size, two exponential decay rates β_1 and β_2, and a stochastic objective function, $f(\theta)$ uses the gradients of the function to update the parameter weights.

This stochastic objective function is also approached in different ways. This function is what defines the loss, and the most common loss function in regression problems is the mean squared error, MSE. When the network predicts a value $\hat{y}$, that corresponds to the true value y, the mean squared error is given as:

$$MSE = \frac{1}{n}\sum_{i=n}^{n}(y_i - \hat{y}_i)^2 \tag{2.23}$$

Generalization and Regularization

The fundamental objective of using machine learning is to achieve generalization. However, when working with finite samples, there is a possibility that the model may fit only the training data and fail to identify a generalizable pattern. This phenomenon is known as overfitting, where the model fits the training data too closely rather than

the underlying distribution. Overfitting can be detected when the data set shows better performance, meaning lower error, on the training data compared to the validation data. To address overfitting, methods that reduce the model's complexity are employed, which are commonly referred to as regularization techniques [48].

There are numerous regularization techniques available to prevent overfitting. Weight decay is a widely used regularization technique that constrains parameter values to a specific range [48]. Another approach that significantly reduces overfitting is dropout [64]. The idea of this approach is to randomly drop units from the neural network during training, which helps prevent units from co-adapting too much.

Feed-Forward Neural Networks

The feed-forward neural network (FFNN) is among the most frequently encountered types of neural networks. These networks include one or multiple "hidden" layers, that can be arranged in a stack, with each one transmitting information to the next layer until the output layer is ultimately reached. A special case of feed-forward neural networks is when all the hidden layers are fully connected. That means that all the neurons in a layer are connected to all the nodes in the following layer. This architecture is often referred to as multilayer perceptron (MLP) [48].

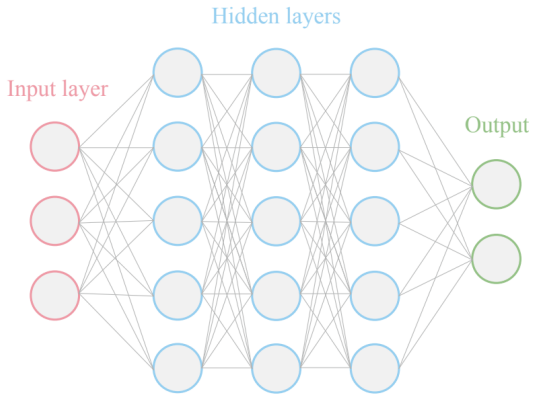

Figure 2.6: An example of a feed-forward neural network with three input nodes, and three hidden layers each consisting of five nodes that are projected into an output of two nodes.

Fig. 2.6 shows an example of a feed-forward neural network with three fully connected hidden layers. The network has three input nodes and two output nodes. Following a fully connected layer, an activation function must be implemented on each hidden unit. These activation functions determine whether a neuron should activate by computing the weighted sum and then adding a bias to it. One commonly applied activation func-

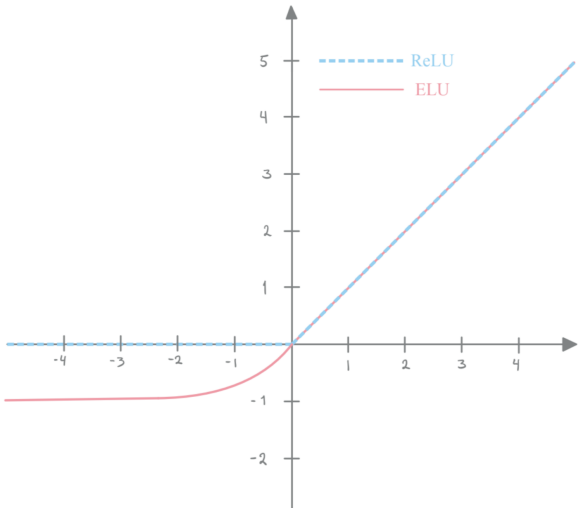

Figure 2.7: A graphic view of the ReLU and the ELU activation functions.

tion is the rectified linear unit (ReLU) activation [65]. This is a very simple approach that returns the maximum of a given element and 0. Another function used to speed up the training process is the exponential linear units (ELUs) [66]. This activation function is a good alternative to ReLU as the ELU function reduces bias shifts by pushing the mean activation closer to zero. The ELU function is given as shown below in equation 2.24, where a is an ELU hyperparameter that is usually set to 1. Figure 2.7 visualizes the plotted functions.

$$f(x) = \begin{pmatrix} x, & x > 0 \\ ae^x, & x \leq 0 \end{pmatrix} \tag{2.24}$$

Previously, these types of networks have been utilized for microstructure modeling and demonstrated promising outcomes [51], [49], [50], [67]. However, it has been limited to less complex models compared to VERDICT, as far as our understanding goes.

2.3.3 Evaluation Metrics

In order to quantify and get a measure of how well the final model is performing, evaluation methods are needed. For regression models, it can be measured either through error metrics or so-called R squared (R^2) metrics. The error metrics measure the difference between the actual target and the predicted value, such as the MSE discussed in section 2.3.2. R^2 is a metric that quantifies the proportion of the variance in the dependent variable y that can be accounted for by the independent variable X. In essence, R^2 measures the degree to which the variance of one variable can be attributed to the variance of another variable. The score is calculated according to formula 2.25 given true value y and predicted value $\hat{y}$ [68].

$$R^2 = 1 - \frac{SS_{RES}}{SS_{TOT}} = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y}_i)^2} \tag{2.25}$$

If a model has an R^2 score of 0.95, it indicates that 95% of the variability in the observed data can be explained by the model's features.

2.3.4 Ethical Considerations of Implementing AI in Health Care

When integrating artificial intelligence (AI) systems into health care, ethical considerations must be taken into account. Ethical issues can arise in different ways [69]. Firstly, they can be epistemic and related to misguided, inconclusive or inscrutable evidence. Secondly, they can be normative and related to unfair outcomes. Finally, they can be related to the traceability of the algorithm. These problems can arise at different levels including individual, interpersonal, group, institutional and societal, or sectoral levels [69]. If actions are not taken in this regard, studies argue that loss of interest in AI could occur due to loss of public trust in the benefits of AI for health care [69].

Chapter 3

Methods

This chapter presents the *in vivo* and simulated MRI data together with the construction of the ANN. Furthermore, the different applications and experiments conducted with the model are presented.

3.1 Datasets

Deep learning algorithms require a large number of annotated data. Even though the approach demonstrated in this paper is self-supervised, there is still a need for annotations in order to validate the quality of the results. To overcome this challenge, several datasets have been simulated based on known ground truth parameters. As a result of this, the research data in this thesis is drawn from two main sources: synthetic data simulation and two sets of *in vivo* MRI data acquired on glioma patients. In the next two sections, these data sets will be described in further detail.

3.1.1 Synthetic Datasets

Ground truth values are challenging to establish in MRI analysis and give difficulties in evaluating the fit of an algorithm. In order to quantify the accuracy of the analyses, synthetic data was acquired by defining the ground truth parameters within plausible ranges, and using combinations of these parameters to form a synthetic data set. In this simulated acquisition process, a combination matrix with seven different parameters from the VERDICT model was used to simulate MRI signals, before noise was added.

Generation of Ground Truth Parameter Array

The VERDICT model described in section 2.2.3 compromises seven parameters to describe the diffusion signals. Only four of these were predicted, as they are specifically relevant in the case of studying the biology of gliomas. These four parameters are the

Table 3.1: Ranges of pseudorandom values used for simulating in silico training data with VERDICT.

Parameter	Minimum value	Maximum Value	Interval
Radius, r	0.02 μm	15 μm	1 μm
f_{IC}	0.1	1	0.1
f_{EES}	0.1	1	0.1
f_{VASC}	0.1	1	0.1
$\lambda_{\|\|}$	$8.0 * 10^{-9} mm * 2/s$	$8.0 * 10^{-9} mm^2/s$	-
d_{EES}	$2 * 10^{-9} mm^2/s$	$2 * 10^{-9} mm^2/s$	-
d_{IC}	$1.2 * 10^{-9} mm^2/s$	$1.2 * 10^{-9} mm^2/s$	-
$\theta(\mu_1)$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	-
$\phi(\mu_2)$	0	0	-

spherical radius of the tumor cell, the volume fraction of the intracellular compartment, the extracellular-extravascular compartment and the vascular compartment. In order to simulate synthetic signals, pseudorandom values were determined for each of these parameters by limiting the values to be restricted within plausible ranges. The remaining three parameters that the VERDICT model is dependent on include three diffusivities for each of the three model compartments. These diffusivities correspond to the intracellular diffusivity which describes water in axons for the spherical compartment, extra-cellular Gaussian diffusivity describes the diffusion in the hindered EES compartment and the stick compartment is described by the vascular diffusivity. These three parameters were fixed to biologically feasible values to reduce the number of unknowns and make the fit computationally more stable. The angles are fixed as the orientation of the vascular compartment does not matter in the modeling, only its contribution. The parameters with their value and range are shown in table 3.1.

Acquisition Scheme

In order to simulate diffusion MRI signals, an acquisition scheme had to be made. This scheme contains all the different acquisition parameters such as gradient direction (x,y,z), gradient strength, diffusion time and gradient duration, which are all used in the simulated acquisition process. Firstly, a large-scale acquisition protocol scheme was constructed to be able to simulate a rich dMRI protocol, to determine whether the method could measure the cell characteristics in such a "perfect" protocol. Although the scanning time required for such a large scheme would be too long, the initial idea behind it was to determine whether this method could be applicable in an optimal setting. After starting with a protocol consisting of about 80 000 measurements, the scheme was narrowed down to a smaller amount of measurements. This was because the large

Table 3.2: Acquisition parameters ranges for simulating in silico training data with VERDICT. Δ is the diffusion time, δ is the gradient duration, G is the gradient strength while $g_{x,y,z}$ is the gradient direction.

Acquisition parameter	Minimum value	Maximum Value	Interval
Δ	10 *ms*	41 *ms*	3 *ms*
δ	10 *ms*	41 *ms*	3 *ms*
G	20 *mT*	80 *mT*	3 *mT*
$g_{x,y,z}$	0	1	-

number of nodes in the linear layer caused memory issues and slowed down the training process. Therefore, a smaller number was chosen based on the available resources. The acquisition scheme was made up of all combinations of G, δ and Δ in plausible ranges that are realistic for a clinical MR system, shown in table 3.2. The TE was set to be constant, as the VERDICT modeling does not take the TE into account. All these combinations resulted in 1155 gradient combinations, in three different directions in addition to one b = 0 image, resulting in a total number of 3466 measurements. This is still considered a very rich protocol, and would not be realistic in an MRI scanning situation.

Data Simulation

The simulated parameter vectors were combined and resulted in 9081 combinations, which were used to simulate a synthetic signal of the same size according to the constructed acquisition scheme. First of all, the VERDICT model was loaded and adapted using the open-source toolbox for diffusion microstructure imaging in Python (DMIPY) [2]. The three different compartments were loaded and given diffusion parameters specific to the case of gliomas. The ball diffusivity, λ_{iso}, the stick diffusivity, $\lambda_{||}$ and the sphere diffusivity, λ_{sphere} are given in table 3.1. As the model was loaded and parameters were set, the signal was simulated by the toolbox implementation of the signal equation described in section 2.2.3. This resulted in simulated MRI values for every 9081 voxels, each with a depth of hardware parameters.

As described in section 2.1.1, Rician noise appears in MRI images. In order to validate and demonstrate, this was tested on the *in vivo* MRI dataset. From figure 3.1, one can see that the voxels in air follow a Rician distribution. As described in section 2.1.1, lower SNR values are correlated to a Rician noise distribution, such as one typically will have at high b-values. This has now been demonstrated in our images, as the air voxels have an SNR of 0.

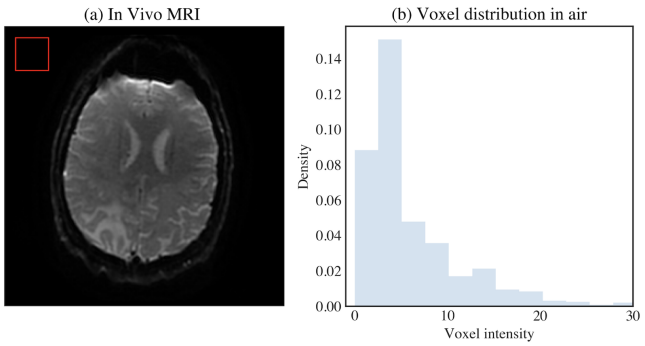

Figure 3.1: Demonstration of Rician noise in air voxels. (a) in vivo MRI with a marked ROI of air voxels. (b) Distribution of air voxels in ROI demonstrating a Rician distribution at low SNR.

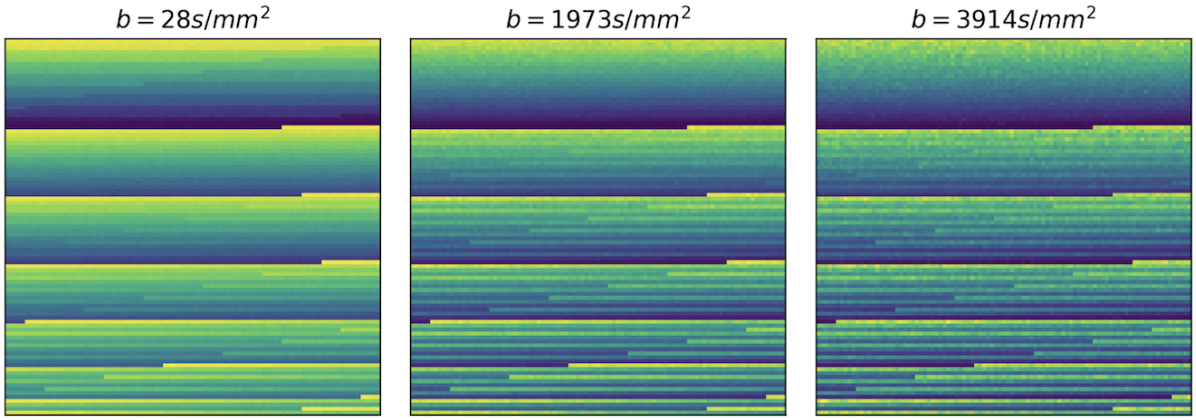

Figure 3.2: In silico training data with increasing amount of Rician noise.

For this reason, Rician noise was added to the simulated signals as well. In order to simulate the correlation between noise and the b-value, the amount of noise was experimented with until reached a formula where signals are fairly recognizable, but with a clear connection to the b-value. By basing the scale on expected SNR for b = 1000 s/mm^2 and b = 3000 s/mm^2 as described in section 2.2, the result was a scaling factor as demonstrated in formula 3.1. Figure 3.2 demonstrates the increasing noise level with the higher b-values, and figure 3.2 shows the noise distribution over different b-values.

$$scale(b) = \frac{10^{-11}}{b[s/m^2]} \tag{3.1}$$

Similar to the approach taken in simulating the training data, a test set was simulated. This set was simulated only using only three different regions consisting of three different ground truth values for each of the four parameters, for easier validation of the

Table 3.3: Range of pseudorandom values used for simulating in silico test data with VERDICT. Only varying parameters are shown, as the remaining parameters are fixed to the same values as shown in table 3.1.

Parameter	Region 1	Region 2	Region 3
Radius, r	5 μm	10 μm	15 μm
f_{IC}	0.2	0.4	0.8
f_{EES}	0.6	0.2	0.1
f_{VASC}	0.2	0.4	0.1

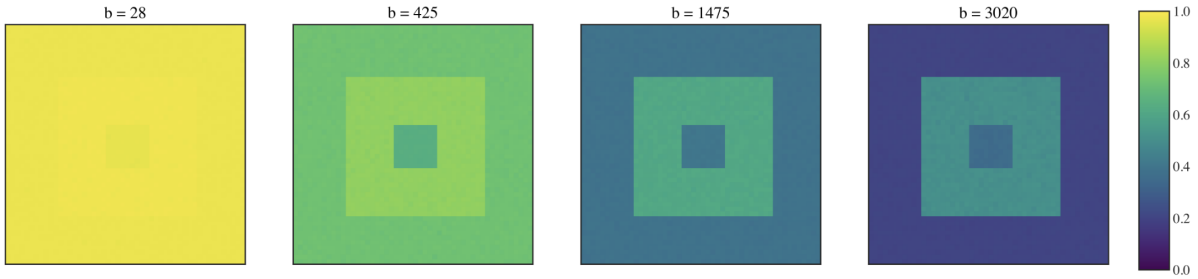

Figure 3.3: The simulated in silico test data with varying b-values.

results by visual inspection. These regions with their ground truth values can be seen in table 3.3, and the simulated data is visualized in figure 3.3

3.1.2 *In vivo* MRI

To test the network under realistic conditions, *in vivo* MRI data that had already been obtained by Zaccagna et al. [43] was utilized. Subjects participating in the study were scheduled for a stereotactic biopsy or resection and gave written informed consent. These MRI examinations were performed using a 3.0 T clinical scanner and a 12-channel head coil (Discovery MR750; GE Healthcare). The gliomas were first located on a T2-weighted image before 16 axial DWI slices were acquired over the region of interest (ROI). The acquisition time for each b-value was 66 s which resulted in a total acquisition time of 330 s abbreviated protocol.

The diffusion-weighted images were acquired with a TR of 2841-3867 ms, and TE of 83.3 - 87.4 ms and with a slice thickness of 5 mm. In this project, data from two patients were used. The first patient was scanned with an abbreviated VERDICT acquisition scheme. This scheme consisted of five b-values, where a separate b = 0 image was acquired to compensate for varying T2-weighted imaging. To assess the robustness of this abbreviated protocol, an extended robust protocol consisting of 40 b-values was undertaken in one patient, within the limitations of the imaging system. The diffusion

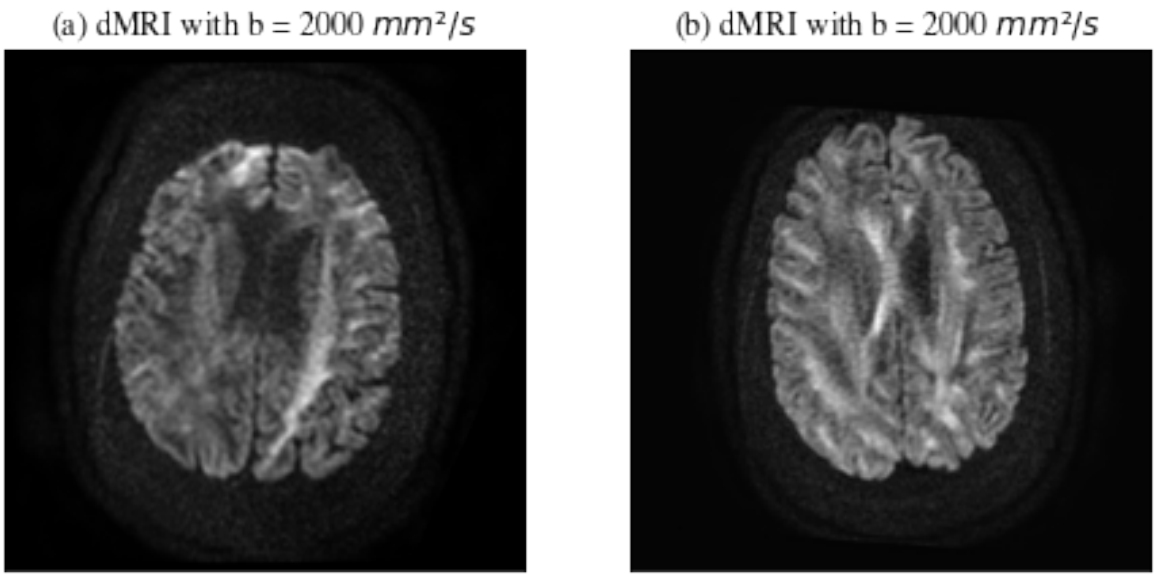

Figure 3.4: Two diffusion-weighted images with b-value 2000 s/mm^2. (a) Shows a slice from the extended protocol with 160 gradient combinations, while (b) shows a slice from the abbreviated protocol with 20 gradient combinations.

parameters for both the abbreviated and robust MRI acquisitions are shown together with a summary of all datasets in table 3.4. Both datasets were corrected for eddy currents and movements using the FMRIB Software Library (FSL) eyes tool to correct for eddy current-induced distortions and subject movements [70]. For usage in the ANN, each b-value image was normalized with its own b = 0 image. Two example slices from the two patients with the different acquisition protocols are shown in figure 3.4.

3.1.3 Simulating Datasets with Less Detailed Acquisition Schemes

Simulations Based on *in vivo* Acquisition Schemes

In order for the trained model to be applicable to *in vivo* MRI data, two additional sets of synthetic datasets were made. The first of these was constructed using the robust acquisition scheme from the *in vivo* data consisting of 40 gradient combinations with three directions and a b = 0 image for each combination. This includes a total number of n = 160 measurements, corresponding to a long acquisition time. The second set was constructed using the abbreviated *in vivo* acquisition scheme with five gradient combinations in three directions and a b = 0 image for each combination. This resulted in n = 20 measurements, which corresponds to a shorter acquisition scheme. Both sets were simulated using the same method as described above, meaning that the same

ground truth parameter array was used, table 3.1. This method also included adding the Rician noise to the images.

Simulation of Explorative Datasets

To explore the capabilities and the limitations of the data, some additional extended datasets were simulated using a variety of acquisition schemes based on a scheme designed for a high-field MR machine. The data were simulated in the same manner as the previously simulated datasets. Firstly, a dataset was simulated using an acquisition scheme from a previously conducted study [12]. This VERDICT study experiments with protocols acquired on a 9.4T MR machine, meaning higher gradient strengths were applied. An acquisition scheme with 46 diffusion weightings was used, having diffusion times Δ = 10, 20, 30 and 40 ms, gradient durations δ = 3 ms for all Δ, and δ = 10 ms for Δ = 30 and 40 ms. Gradient strength, G, varied from 40 to 400 mT/m in 10 steps of 40 mT/m for δ = 3 milliseconds and |G| = 40, 80, and 120 mT/m for δ = 10 milliseconds. This setup results in 46 different gradient combinations in three different directions, with n = 145 measurements in total.

This acquisition protocol was extended in several ways, with the first extension including the same amount of measurements as the original scheme, but with extended ranges of Δ and larger steps between values. This made a maximum b-value of $13454 s/mm^2$. The second extension included higher values of δ and a higher range of Δ. This made a maximum b-value of $18549 s/mm^2$ and a total number of n = 172 measurements. The third extension was making the steps in the diffusion times even smaller, with only 5 ms between values of diffusion times, Δ, which resulted in the same maximum b-value, but with a total number of n = 436 measurements.

As demonstrated in section 3.1, the amount of Rician noise increases with increasing b-value, which means that this network should be even more robust towards noise than for lower b-value acquisitions. Even though it is interesting to see whether high b-values could provide accurate results in this model, it is necessary to test them with a high amount of noise. The added noise was scaled by an SNR of $1e^{11}/b$ which resulted in an SNR ranging between 3 for the highest b-value which results in nearly unrecognizable signals, and 100 for the lowest b-value.

3.1.4 A Summary of All Datasets

Throughout the project, a variety of datasets, consisting of both simulated and *in vivo* data, were utilized to explore the method's potential. Table 3.4 provides a comprehensive summary of all the datasets, including their details.

Table 3.4: Overview of all used datasets. n is the total number of measurements in the dataset, and acquisition parameters include diffusion time, Δ, δ is the pulse duration and G is the gradient strength. [1] Acquisition scheme from the in vivo data from Zaccagna et. al [43], [2] acquisition scheme from Panagiotaki et al. [12].

	Dataset	**n**	**size**	**Acquisition scheme values**		
				Δ [ms]	δ [ms]	G [mT]
in silico	**Large-scale train set**	3466	[9180,3466]	{10,13,...40}	{10,13,...40}	{10,13,...80}
	Large-scale test set	3466	[2500,3466]	{10,13,...40}	{10,13,...40}	{10,13,...80}
	Abbreviated[1]	20	[256,256, 252,20]	{24,32,43}	{5,14,26}	{30,44,76}
	Robust[1]	160	[256,256, 16,160]	{23,32,43}	{10,12,...,25}	{10,20,...60}
	Extended (high gradient strengths[2])	145	[9180,145]	{10,20,30,40}	{3,10}	{40,80,...,400}
		145	[9180,145]	{10,25,...,55}	{5,12}	{40,80,...,400}
		172	[9180,172]	{10,20,...,50}	{8,15}	{40,80,...,400}
		436	[9180,436]	{10,15,...,60}	{5,12}	{40,80,...,400}
in vivo	**Robust**[1] **MRI**	160	[256,256, 16,160]	{23,32,43}	{10,12,...,25}	{10,20,...60}
	Abbreviated MRI[1]	20	[256,256, 252,20]	{24,32,43}	{5,14,26}	{30,44,76}

3.2 Tissue Parameter Fitting with Neural Network

Once the data was constructed and processed, it was employed to train an artificial neural network (ANN) with the objective of estimating the tissue parameters considered by the VERDICT model. The network is inspired by the less complex implementation of a DL approach to fitting a simple ball-stick model done by Lim et al. [49]. The following sections will describe the steps in the ANN training process.

3.2.1 Artificial Neural Network Architecture

As the VERDICT model is a voxel-wise model, the ANN works on the image data voxel-wise. This means that the input layer consists of signals from the same voxel, all with different weightings according to its b-value. With the large-scale simulated dataset, this results in an input layer with the tensor size 3466, i.e. the number of measurements. These voxel signals are then fed forward through the MLP architecture described in section 2.3, meaning three fully connected layers followed by an ELU activation function, each with input and output size of the number of measurements, n.

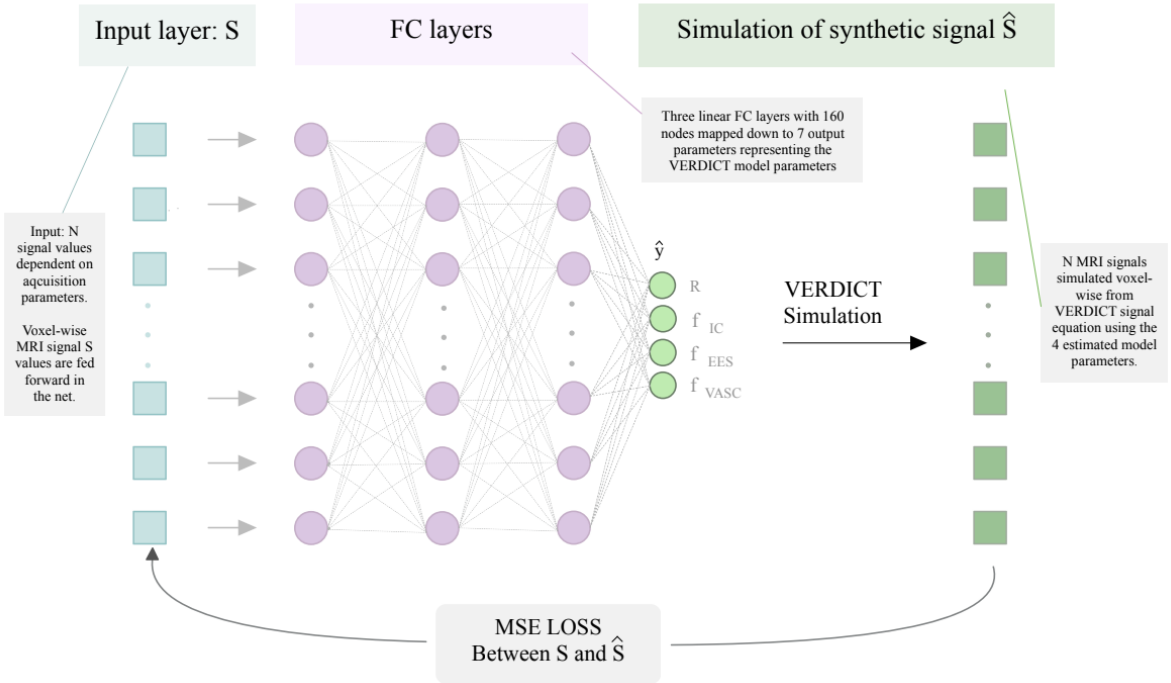

Figure 3.5: Artificial neural network architecture. An input layer which is the voxel-wise MRI signal (S), is forwarded through three fully connected (FC) layers. These are mapped down to four output values which consist of the radius (R), and the 3 volume fractions: The intra-cellular (f_{sphere}), The extracellular-extravascular (f_{ball}) and the vascular f_{vasc} compartment. These were then used to simulate a synthetic MRI signal, $\hat{S}$. The mean squared error (MSE) loss is then computed.

After being fed forward through the hidden layers, the signals are mapped down to four values, representing cell radius, R, and the three volume fractions, f_{IC}, f_{EES} and f_{VASC}. The remaining parameters that the VERDICT model is dependent on are fixed to the same values as used when the data was constructed. From the four free parameters and the fixed parameters, a new set of MRI signal values from the voxel can be estimated using the VERDICT model equation. With this new signal, a loss can be computed. This is done by computing the MSE loss between the synthetic and the original signal. With this loss, a backward pass is done and the weights in the fully connected layers are updated. This means that the tissue parameters can be collected directly from the training process itself. The range of fitted parameter values is not constrained by the network. A schematic representation of the architecture can be seen in figure 3.5.

VERDICT Simulation

A very central part of the network is the simulation of the new MRI signal, and a further explanation is needed. As the previously demonstrated work uses a simple IVIM model which only compromises a ball and a stick compartment is a quite simplistic model [49], the signal simulation is a more complex process in this work. The signal values from the voxel are mapped down to the four free model parameters, which are all used to simulate signals from either the IC, EES or vascular compartment. The three fractions f_{IC}, f_{EES} and f_{VASC} are then multiplied with the respective compartments' signal to compute the total signal value. All of these three compartments are modeled according to their specific signal equation.

The EES compartment which is modeled by a diffusion tensor is in this case constrained to be isotropic because of the cell's isotropic arrangement and is therefore represented as a ball. The ball compartment has the most simple signal equation as described in section 2.2.3, and is only dependent on the b-values and d_{EES}, which is in this case set to be constant. This means that the signal from this compartment does not change, although the fraction of it in the total signal is determined. The EES diffusion constant, d_{EES} is set to $2 * 10^{-9} m^2/s$.

The vascular compartment model is also modeled as a diffusion tensor, but in this case, we assume the compartment to have a high degree of anisotropy, which means that we orientate it in one single direction, explained by the angle μ. This angle, mu, is again a combination of the two angles, θ and ϕ. While these are passed through the network,

they are calculated in 1D unit sphere coordinates. These are then converted to Cartesian coordinates before the dot product is computed in the signal equation. However, as we assume anisotropic diffusion along one single direction, these angles are fixed. The vascular diffusion constant, $\lambda_{||}$, is set to $8 * 10^{-9} m^2/s$, while θ and ϕ is set to $\frac{\pi}{2}$ and 0, respectively, as these were assumed isotropically orientated.

The signal from the sphere compartment is more complex, and thus a more computationally expensive process. This implementation is derived from the DMIPY Github [71] way of computing the spherical component as described in section 2.2.3. Presented in DMIPYs source code [71] are the first 100 roots of the equation, which is also used in this project's calculations. Furthermore, the implementation has been adjusted to take in not only one value of δ, Δ and gradient directions at a time but a vector of these. This was done to facilitate our new method of doing the modeling, which differs from the way it has been done in DMIPY.

Now, these three signals are combined using the fraction parameters, which decide the proportion of which each compartment of the specific voxel consists.

Implementation

The network was implemented on Python 3.9.12 [72] with PyTorch version 1.13.0 [73]. During the introductory testing stage, models were trained on a MacBook 3.1 GHz Dual-Core Intel Core i5. All final models were eventually trained on an AMD Ryzen 1950X 16-core CPU and GeForce GTX 1080 Ti GPU with Cuda version 12.0.

3.2.2 Model Training on a Simulated Large-scale Dataset

The initial model training was done on the large-scale simulated dataset (n = 3466). This would not be applicable in a clinical situation as this would require a very long scanning time, but it has the potential to demonstrate whether or not this modeling technique can reveal detailed information about the tissue. The data was loaded into the developed framework using a PyTorch DataLoader [73] before it was forwarded in batches through the neural network. The model was initialized with model parameters similar to the ones that gave the best results in the paper of J. Lim [49]. This included a learning rate of 0.001, 100 epochs and a dropout rate of 0.5. However, in order to improve the results, the model was trained with many different combinations of both model parameters and processed input data such as varying noise levels and different

ranges of radius. Also based on the results of the same paper, the ADAM optimizer and the MSE loss function described in section 2.3 is used for all training.

The training was done with the simulated data with parameters described in table 3.1, but the radius was experimented with. The initial training was firstly done with a maximum radius of 30 μm, but since this value is outside the expected ranges of a glioma cell radius, the model was trained on simulated data with lower radius ranges with a maximum of 20 and 15 μm. This was done to investigate if training on data with lower radius values could improve the outcomes from model training.

At the early stage of training, the model's MSE loss was continuously monitored to manage how well the parameters were estimated during training. Even though training for more epochs might provide a better fit, it is also more time-consuming. Therefore, in order to deal with the compromise between accuracy and time consumption, we chose to proceed with the number of epochs where the loss decrease seems to have stagnated. To find this point, the general trend of MSE losses was visualized across multiple models, aiming to identify the minimum number of epochs required to achieve a converging loss. This approach helped in reducing the number of epochs needed to train an accurate network without it being too slow.

Test Metrics

To assess the performance of the models on the training data, R^2-scores were calculated from the ground truth parameter vectors and the predictions on training data. This was done using the Scikit-Learns [74] build-in method that calculates the R^2-score as described in section 2.3.2. Additionally, prediction maps and scatter plots were visualized. As a curiosity and for later comparison, calculations were made of the mean, standard deviation and mean squared errors of the predictions at one expected value in each of the four parameters. These four example values were chosen based on the previously conducted studies on glioma tissue values described in section 2.2.5. The statistical metrics were therefore calculated on predictions of radius at ground truth value 8 μm, and IC, EES and vascular volume fractions 0.2, 0.7 and 0.1, respectively.

3.2.3 Predicting on an Unseen Simulated Large-scale Dataset

Even though this method would require retraining to each specific acquisition protocol, it is of interest to see how well the model generalizes. If the generalization is good, it

shows the possibility to be trained for different acquisition schemes and then applied right away with a very low processing time. To discover if this was possible, the model trained on the large-scale simulated *training* dataset (n = 3466) was applied to the simulated large-scale *test* set (n = 3466).

To generate predictions with an already trained model, the model and dataset were loaded in batches using the PyTorch DataLoader before the trained weights were applied to the data and predictions were made. Predictions from the three regions of ground truth values were analyzed. To assess the model's performance on the test data, R^2-scores, parameter maps, and scatter plots were calculated as done on the training data.

3.2.4 Model Training on Simulated Robust and Abbreviated Protocol

After investigating the training results and generalizability from training and testing on the large-scale dataset (n = 3466), model fitting was carried out using the abbreviated and robust simulated data (n = 160 and n = 20). These synthetic datasets described in section 3.1 were both composed of acquisition schemes used for *in vivo* MRI scans, making the trained models applicable to the *in vivo* datasets. These were first trained using the model that included noise in the data, as the purpose was to fit noisy data. Next, they were also trained on data without noise, to investigate whether or not such shallow acquisition hypothetically could provide enough information to perform good predictions.

Similarly to the previous model training, R^2-scores were computed on the training predictions based on the ground truth parameter array, table 3.1, used when simulating signals. From the same ground truth values, scatter plots were made to visualize the prediction error. Additionally, statistical measures of the MSE error in the same four example values (R = 8 $\mu m, f_{ic} = 0.2, f_{EES} = 0.7, f_{VASC} = 0.1$) used in the large-scale dataset evaluation. This was done in order to compare the abbreviated and robust fitting with the fitting of the large-scale dataset.

After the model had been trained on these datasets, the trained weights were then applied to the *in vivo* data. Together with the data, a brain mask was loaded using the PyTorch DataLoader [73], and the predictions of voxels outside the brain mask area

were set to zero as these voxels only contain air. The brain masks were generated using MDTs built-in method [15].

3.2.5 Model training on Extended protocols

In order to find the balance between the unrealistic scenario presented by the large-scale dataset (n = 3466) and the potentially insufficient information provided by the abbreviated (n = 20) and robust (n = 160) protocol, it is desirable to discover a method that falls somewhere in the middle. As a result, additional acquisition schemes based on MR acquisitions designed for a high-field MR machine (9.4 T) have been investigated. The model was trained on the four experimental extended datasets with high gradient strengths (n = 145, 145, 172, 436) described in section 3.1.3. Predictions from training were compared to the ground truth parameter array by computing R^2-scores and scatter plots were made. Also for these model fittings, statistical measures of the MSE error were calculated in the same manner as for the previous model fittings.

3.3 Comparison with Existing Methods

After tuning the network and demonstrating the suitability of the developed ANN, its performance was compared to two existing methods described in section 2.2.4. These were firstly compared in terms of computational time before the accuracy of these methods was explored on data containing ground truth values. All approaches were tested with the AMD Ryzen 1950X 16-core CPU or GeForce GTX 1080 Ti GPU with Cuda version 12.0.

3.3.1 Comparison of Computational Time

For all three methods, the time consumption per voxel was calculated. This amount of voxels included all data points from the *in silico* data and all voxels within the brain mask for the *in vivo* data.

ANN

The comparison took into account the fact that the proposed artificial neural network (ANN) needs to be trained for each specific acquisition scheme. Therefore, the ANN is compared both in terms of training the network and in terms of applying an already trained network to a new scan. Once trained for a specific protocol scheme, a generalizing model can potentially be directly applied without requiring any further training. The time consumption of the ANN was tested on the large-scale (n = 3466) *in silico* dataset, and both the robust (n = 160) and abbreviated (n = 20) *in vivo* protocol.

DMIPY

The first method that the ANN is compared to is the DMIPY framework that was described in section 2.2.3. To ensure that the approaches have the same origin, all diffusion parameters except volume fractions and radius are fixed using DMIPYs built-in function for fixing parameters. The parameters are fixed to the same values as in the ANN, described in table 3.1. The MIX optimizer as described in section 2.2.4 is used as it is used in the DMIPY code example for fitting the VERDICT model, and it was described to be efficient in finding the global minimum in models with many compartments as we have here [71]. The time consumption of the ANN was tested on the large-scale (n = 3466) *in silico* and robust *in vivo* protocol (n = 160). The abbreviated *in vivo* protocol (n = 20) was excluded because it has significantly more voxels than the longer protocol, and the time per voxel was expected to be nearly identical. The

DMIPY toolbox was implemented with DMIPY version 1.0.4 [71], Dipy 1.5.0 [75] and Scipy 1.7.3 [76]. The AMD Ryzen 1950X 16-core CPU was used.

MDT

The MDT framework with version 1.2.7 [15], described in 2.2.3, constitutes the second method. In order to apply this approach it was necessary to obtain a brain mask, which could be generated from the toolbox itself. To convert the acquisition scheme, all the relevant scheme parameters were extracted and utilized to create the required protocol file through the toolbox. The fitting process was first conducted with standard settings, which utilized the Powell optimizer. Even though this method had been described to outperform the two others, the two other methods were also tested. This method was also tested on all the same datasets as the ANN. The GeForce GTX 1080 Ti GPU with CUDA VERSION 12.0 was used.

3.3.2 Comparison of Accuracy

To further investigate the existing methods, results obtained by fitting these frameworks are analyzed. The results obtained during the computational time comparison on the large-scale *in silico* dataset (n = 3466) were visualized to investigate the fit of the two existing methods, in order to compare with the developed ANN. Results from fitting on the *in vivo* data were not analyzed, as these do not contain any ground truth values to validate the accuracy of the fit. Since the MDT toolbox requires a mask together with the data, a binary array with the same shape as the *in silico* data consisting of ones were made and applied when fitting with the MDT.

Chapter 4

Results

This chapter presents results and statistical testing from ANN implementations on MRI data. Additionally, results from comparisons with existing methods are presented.

4.1 Artificial Neural Network Training Results

Results from training on the large-scale simulated dataset (n = 3466) are presented with the predicted parameter maps from the training process. After this, training was done on the simulated abbreviated (n = 20) and robust (n = 160) protocols, before these trained models were applied to *in vivo* MRI data. Furthermore, results from training with the extended acquisition schemes (n = 145, 145, 172, 436) are presented.

4.1.1 Model Training and Selection on Large-scale Dataset

The ANN was trained with different settings as described in section 3.2. After deciding which input data to include, the ANN was trained with different model hyperparameters. All training processes that gave decent results with their outcomes are shown in table 4.1. From the two learning rates, it was clear that a learning rate of 0.0001 gave the best results. It is also noticeable that the training loss was generally higher when adding noise to the simulated signal. Highlighted are the model training with the best loss and highest training accuracy both with and without added noise.

Table 4.1: Overview of the training process. Training parameters are learning rate (LR), weight decay (WD), and dropout which is the applied dropout rate.

Model	Epochs	LR	WD	Dropout	Noise	Best loss
Model0	100	0.0001	0	0.5	×	0.006
Model1	50	0.0005	$1e^{-5}$	0.5	×	0.010
Model2	20	0.0005	0	0.5	×	0.012
Model3	200	0.0001	0	0.5	×	0.014
Model4	20	0.0005	0	0.5	✓	0.091
Model5	200	0.0005	0	0.5	✓	0.043
Model6	200	0.0001	0	0.5	✓	0.010

It was observed that the model fitting did not perform equally well on data that were simulated using a wide range of radius values. The training performance was remarkably better when fitting radius values between 0.02 and 15 μm rather than the firstly proposed range between 0.02 and 30 μm. As glioma cell radius larger than 15 μm are less realistic *in vivo*, it was decided to proceed with the training process based on simulated data with a smaller range of radius values, meaning between 0.02 and 15 μm.

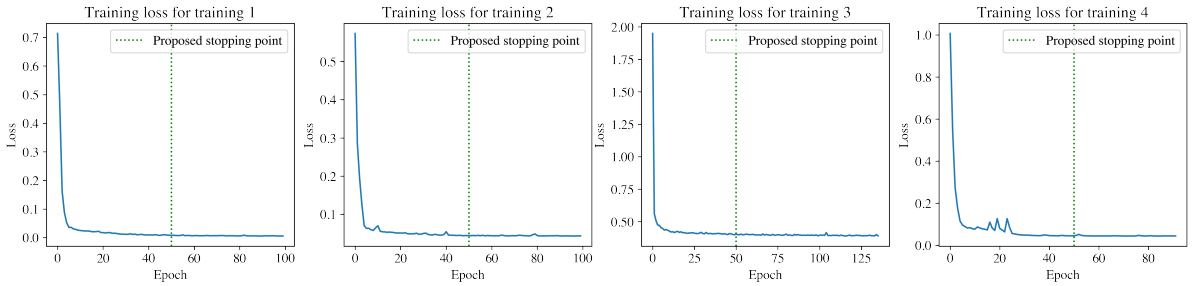

Figure 4.1: The loss values plotted for several different training processes. The blue line shows loss values, while the vertical green line shows the proposed stopping point at 50 epochs.

As seen from figure 4.1, all the training losses follow a pattern where it decreases for each epoch. As this was at an experimental stage, all training was done with different amount epochs, and the axis of all plots, therefore, does not match. However, from analyzing these plots, the number of epochs was then proposed to a maximum of 50 epochs as seen from figure 4.1, which is where the graphs stop showing significant growth or change for all of the trained models. As a result of this analysis, the models to proceed with were the highlited models from table 4.1, but training can be stopped at 50 epochs.

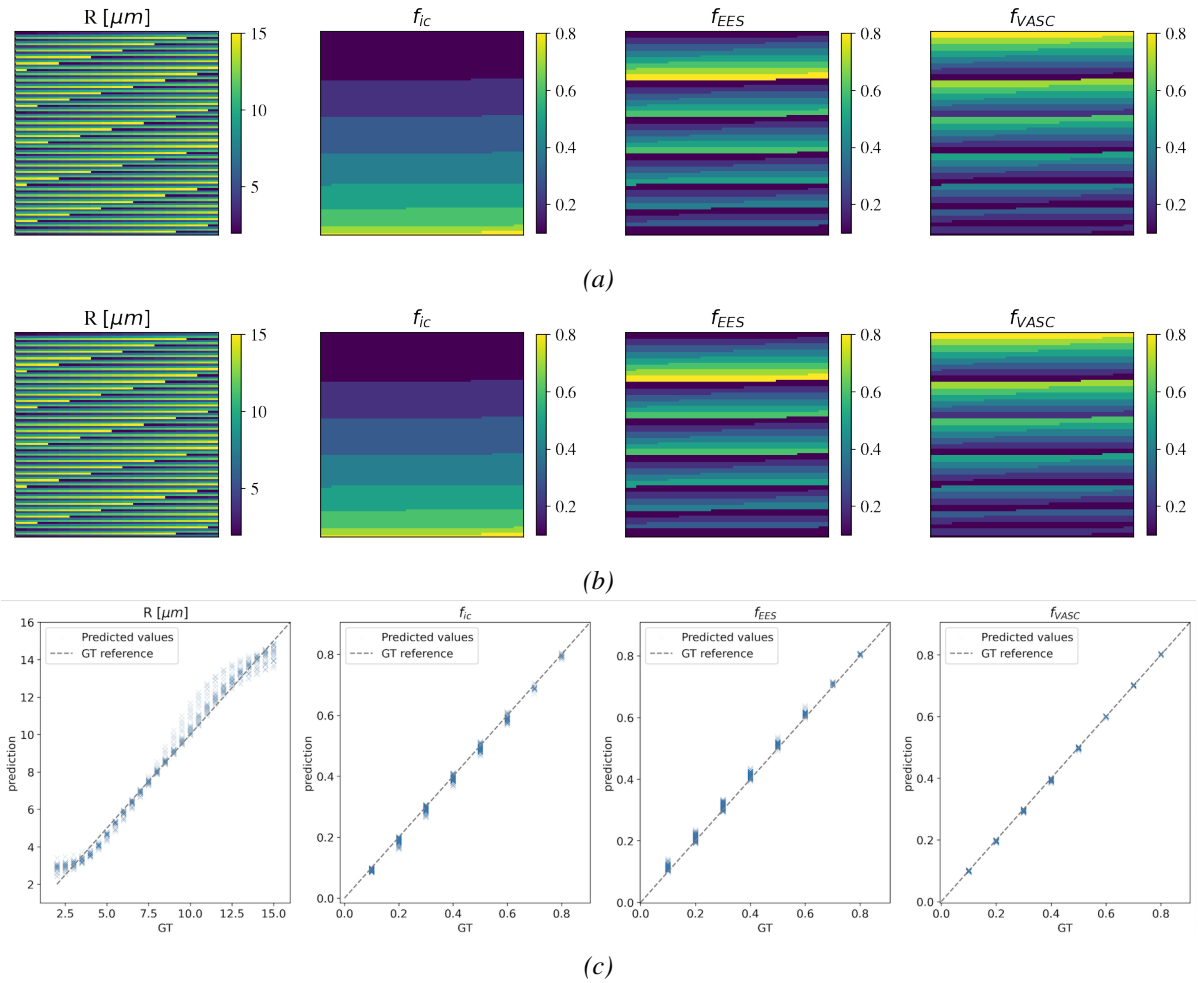

Figure 4.2: Comparison of ground truth (GT) and predicted parameter maps on the large-scale dataset. Showing cell radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}). (a) Ground truth parameter maps. (b) The predicted values from the constructed neural network. (c) Scatter plots over the ground truth values on the x-axis and predictions on the y-axis.

Maps from training the final model were visualized together with the ground truth maps to validate its strong alignment, as affirmed in figure 4.2 (a,b). The results were also visualized in scatter plots, see Fig. 4.2 (c), which shows that the predicted values do not deviate particularly from the ground truth values as they follow a linear path closely. However, one can see that the deviations are typically larger for lower values of the volume fractions, while for the radius, the deviations are higher in both the low and high ranges. To support the plots and visual results, R^2-scores were calculated for all four parameters, table 4.2. In this table, scores from the model without noise are also included. All R^2-scores are higher on the model trained on data without noise. Visualization of results from training on data without noise was not done, as an infinite SNR is not achievable.

Table 4.2: Metrics obtained from training the large-scale protocol with and without noise. R^2 scores on the predicted radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}).

Data	R^2-scores			
	Radius (R)	f_{ic}	f_{EES}	f_{VASC}
Data without noise	0.982	0.995	0.994	0.999
Data with noise	0.831	0.980	0.970	0.996

Additionally, statistical measures of mean, standard deviation and MSE in four values expected to find in glioma tissue are shown in table 4.3. These values support the previous finding that the most accurate predicted parameter is the vascular volume fraction, and they also show that the predictions have small errors.

Table 4.3: Statistical measures on example values (R = 8 $\mu m, f_{ic} = 0.2, f_{EES} = 0.7, f_{VASC} = 0.1$) on artificial neural network training predictions from the large-scale dataset. Mean, standard deviation and mean squared error (MSE) on the predicted radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}).

	Radius [μm]	f_{IC}	f_{EES}	f_{VASC}
Mean	8.04 ± 0.15	0.187 ± 0.009	0.710 +-0.004	0.098 ± 0.002
MSE	0.02	0.0002	0.0001	6×10^{-6}

4.1.2 Model Prediction on Large-Scale *in silico* Data

As the extracted maps from the training process had shown to be close to ground truth values, the two final models were also tested on the large-scale simulated test set (n = 3466). Since weights were already trained, the network predicted the same predictions for all pixel values within the same region. The results from the testing along with the ground truth values are shown in figure 4.3 (a-b), where it can be visually seen that the predictions match the ground truth. Furthermore, Fig. 4.3 (c) shows the predicted values plotted alongside the ground truth reference, with very limited deviations. Although this set contains only three different ground truth parameters to predict, one can see from the scatter plots that the predictions follow the same trend as in the training data. The vascular component, f_{VASC} has the most accurate prediction, where the EES compartment has been slightly overpredicted and the IC compartment slightly underpredicted.

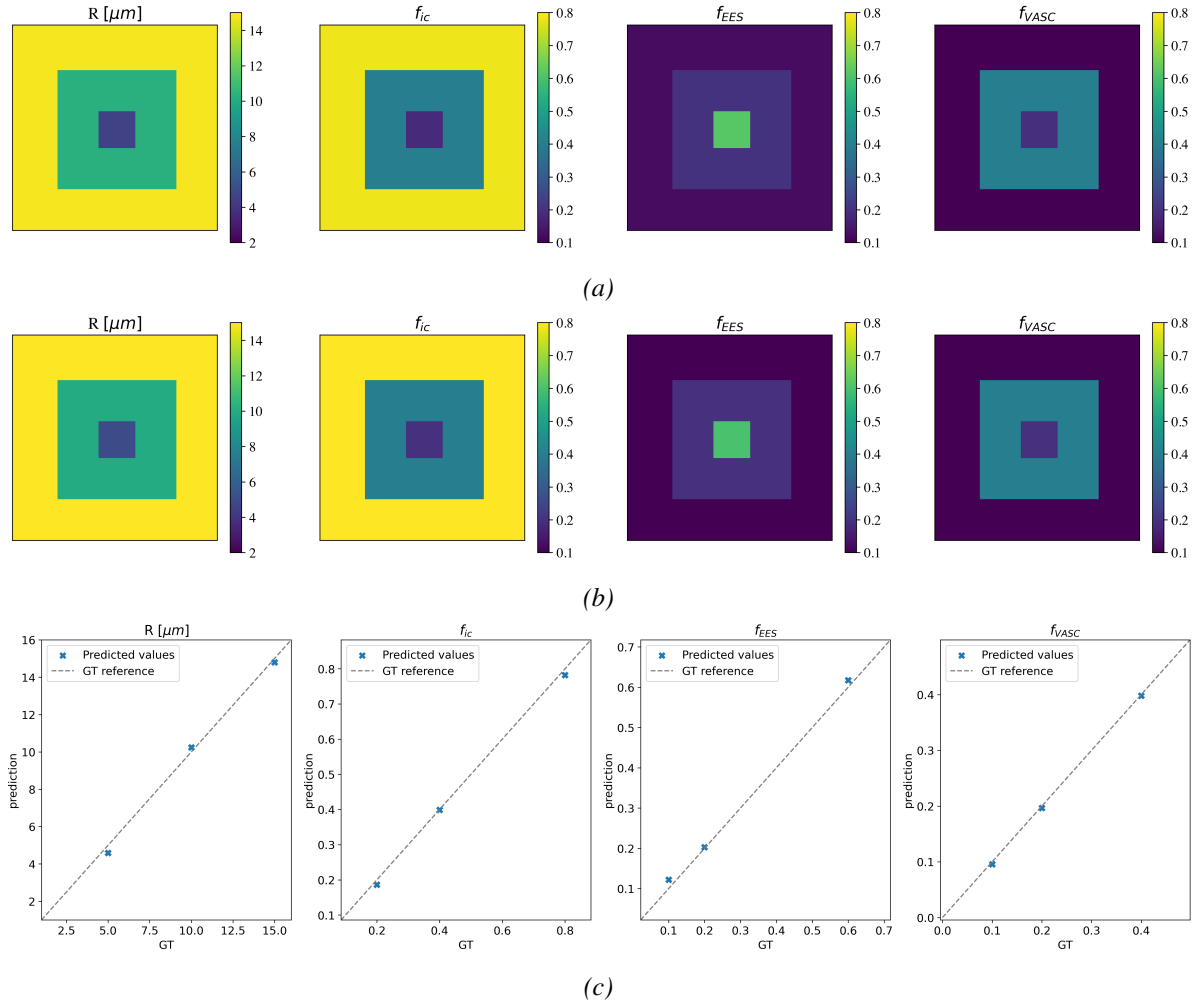

Figure 4.3: Comparison of ground truth (GT) and predicted parameter maps on the test set. Showing radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}). (a) Ground truth parameter maps. (b) The predicted values from the constructed neural network. (c) Scatter plots showing the GT values on the x-axis and predictions on the y-axis, together with a GT reference line.

The R^2-scores, table 4.4, support the visually shown results on the test set. The radius predictions perform even higher on the test set, IC and vascular volume fractions perform almost the same. The EES volume fraction has a slightly lower score.

Table 4.4: R^2 scores from the test data. Calculated from the predicted radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}), and vascular volume fraction (f_{VASC}) on the test dataset.

Data	R^2-scores			
	R [μm]	f_{ic}	f_{EES}	f_{VASC}
Data without noise	0.993	0.995	0.962	0.999
Data with noise	0.576	0.999	0.986	0.995

4.1.3 Training on Simulated Robust and Abbreviated Protocols

To assess the performance of the trained model on the simulated robust (n = 160) and abbreviated (n = 20) protocol, scatter plots showing the relationship between the ground truth and predictions in training data were generated for these datasets as well, Fig. 4.4. The results are dramatically less accurate than for the large-scale dataset (n = 3466), and predictions do not follow the linear relationship to the same degree. In the robust protocol, radius predictions do not manage to predict values in the low or high range, and there are big deviations from the ground truth. Predictions on the abbreviated protocol show that radius values are predicted to a nearly constant minimum value. Table 4.5 supports these visual results, with in general very low R^2-scores compared to training on the large-scale data (n = 3466). The R^2-scores seen from the table also indicate that training the model in the hypothetic situation without any noise actually made the accuracy worse than training with added noise.

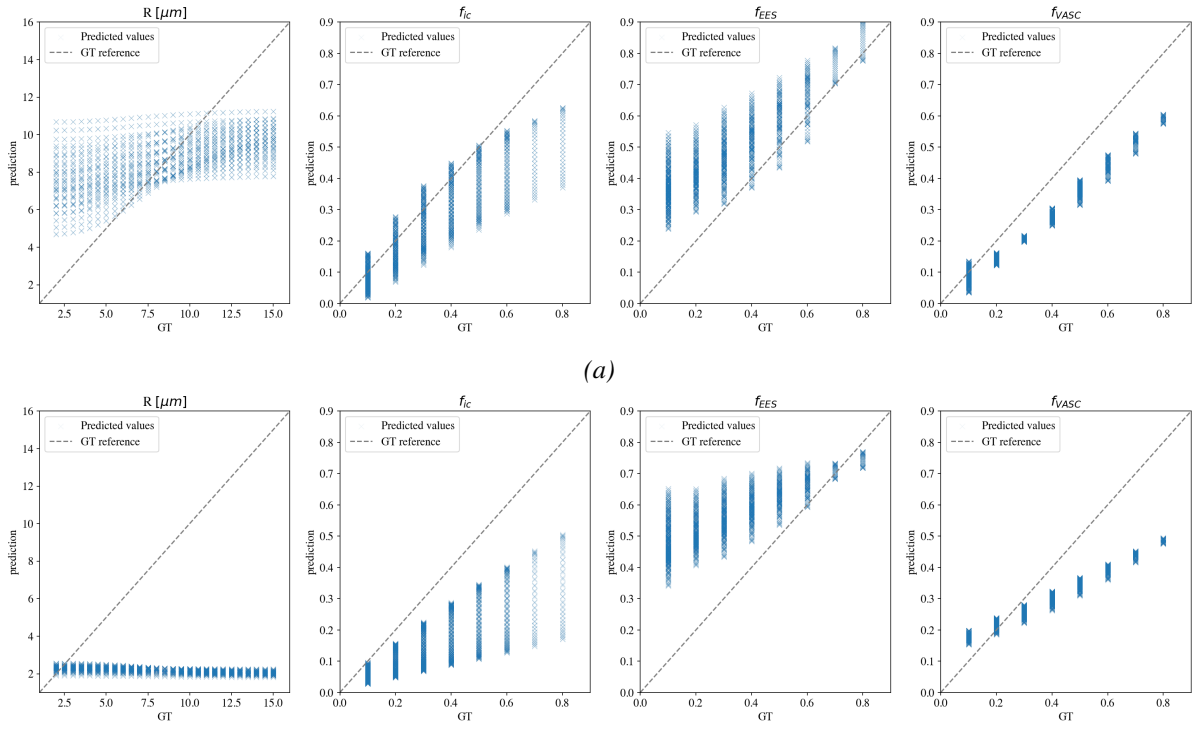

Figure 4.4: Scatter plots of ground truth (GT) in training data and predicted parameter maps. Showing cell radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}). (a) The model trained on the robust protocol, (b) the model trained on the abbreviated protocol.

Table 4.5: Training results from training on the realistically constructed data. R^2 scores on the predicted radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}).

Data	Noise	Best loss	Training time [s]	R^2-scores			
				$R[\mu m]$	f_{ic}	f_{EES}	f_{VASC}
Robust protocol	Yes	0.41	693	-0.301	0.606	-0.069	0.670
	No	0.42	819	-0.450	0.588	-0.082	0.706
Abbreviated protocol	Yes	1.59	669	-1.401	-0.290	-1.048	0.556
	No	1.62	630	-1.53	-0.401	-1.494	0.420

For comparison with the large-scale fitting, table 4.6 shows the calculated mean, standard deviation and MSE on four example values (R = 8 $\mu m, f_{ic} = 0.2, f_{EES} = 0.7, f_{VASC} = 0.1$). The MSE was also compared to the calculated MSE on the large-scale model training. This shows that the error is at the minimum 4 times bigger in the robust protocol (n = 160) fitting compared to the large-scale protocol (n = 3466) fitting. On the other side of the scale, the radius error is 1705 times as big in the abbreviated protocol (n = 20) fitting.

Table 4.6: Statistical measures of example values (R = 8 $\mu m, f_{ic} = 0.2, f_{EES} = 0.7, f_{VASC} = 0.1$) on artificial neural network training predictions on the abbreviated and robust dataset. Mean, standard deviation and mean squared error (MSE) on the predicted radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}). Relative MSE error is the calculated MSE compared to MSE on large-scale predictions.

		R $[\mu m]$	f_{IC}	f_{EES}	f_{VASC}
Robust protocol	Mean	8.64 ± 0.99	0.167 ± 0.055	0.709 ± 0.018	0.089 ± 0.027
	MSE	1.392	0.0040	0.0004	0.0009
	MSE/MSE large-scale	69.6	20.0	4.0	150.0
Abbreviated protocol	Mean	2.17 ± 0.12	0.101 ± 0.036	0.762 ± 0.040	0.181 ± 0.011
	MSE	34.1	0.011	0.005	0.007
	MSE/MSE large-scale	1705.0	55	60.0	1166.7

4.1.4 Predictions on *in vivo* MRI data

As it was shown that the model was capable of making predictions on an unseen large-scale test set (n = 3466), the approach was tested on *in vivo* data using the models trained on the robust and abbreviated protocol. Fig. 4.5 (a) shows parameter predictions on the robust *in vivo* MRI protocol and fig. 4.5 (b) shows predictions on the abbreviate protocol. The predictions of radius for the abbreviated protocol exhibit un-

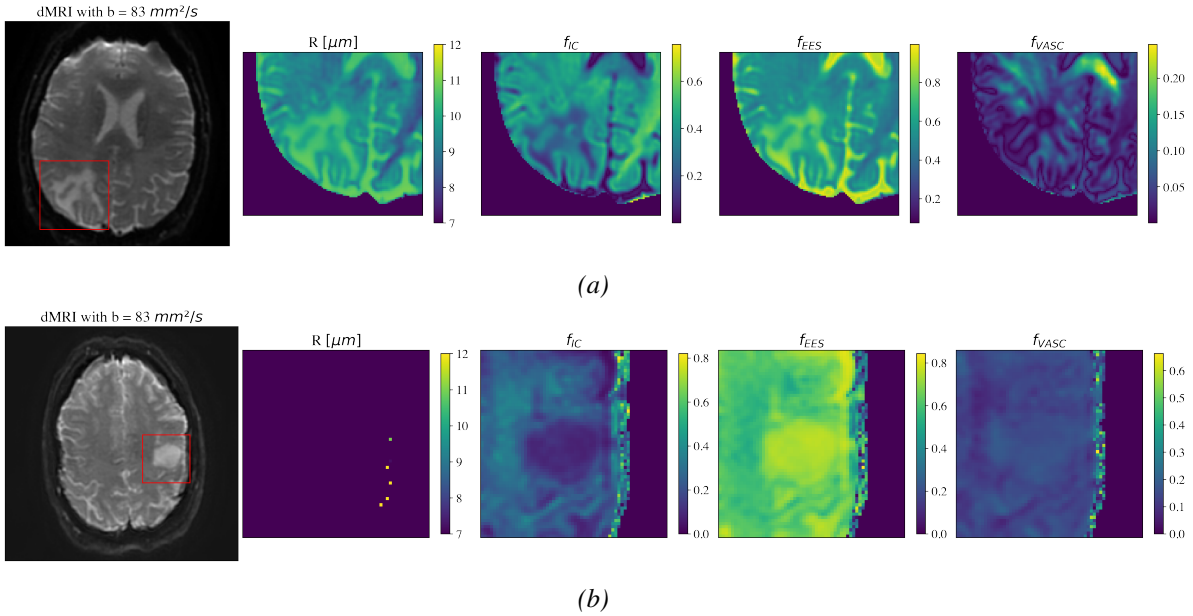

Figure 4.5: in vivo diffusion MRI with VERDICT parameter predictions. The ROI is highlighted in two dMRI scans, and parameter predictions are shown for the corresponding region. The parameters include the radius and three volume fractions: intracellular (f_{IC}), extracellular-extravascular (f_{EES}) and vascular (f_{VASC}) in an MRI protocol consisting of 160 gradient combinations.

realistically low and consistent values, whereas the other parameter predictions demonstrate the ability to capture the tissue structure.

The abbreviated protocol produces a range of seemingly arbitrary values in the edges of the brain, which can be explained by the mask not fitting perfectly on the brain tissue. The model generates values in the skull, which is beyond the scope of the VERDICT model and hence not reflective of actual values, and should thus be disregarded. Similarly, predicted values in the cerebrospinal fluid should also be disregarded.

4.1.5 Predictions on Extended Protocols

Figure 4.6 shows the results from training with different gradient combinations from the extended acquisition schemes (n = 145, 145, 172, 436). (a) shows the parameter predictions on the new acquisition scheme with 145 measurements, while (b) shows the predictions on the extended acquisition scheme with higher b-values but the same number of measurements. (c) Shows the extended version with both more measurements and higher b-values compared to (a). Furthermore, (d) shows predictions from training with an acquisition scheme with both more measurements and higher b-values. Table 4.7 explains the details of the four different acquisition schemes. Training time for all models was under 700 s.

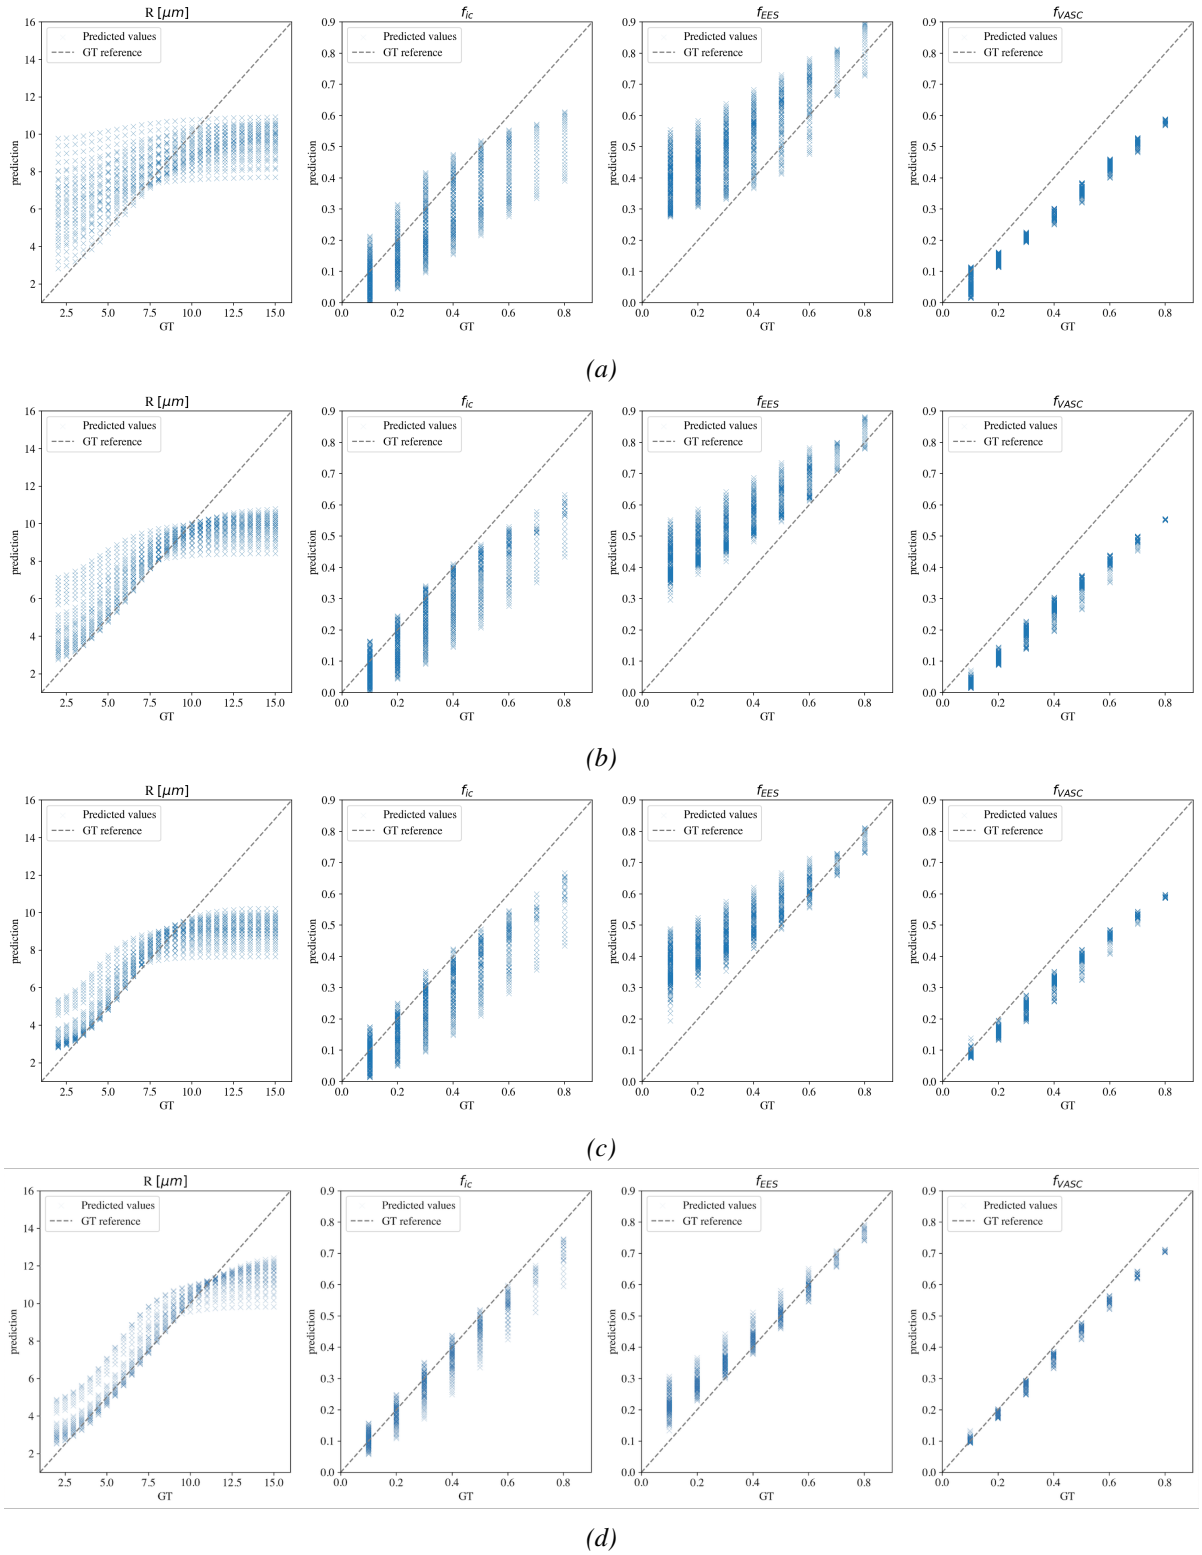

Figure 4.6: Comparison of ground truth (GT) and predicted parameter maps on the training set for the extended acquisition schemes. Showing radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}). (a) Predictions from a real acquisition scheme with a total of 145 measurements, (b) scheme extended with larger but less dense δ-range and higher Δ, (c) extension to higher b-values with 172 measurements, (d) a further extension of the acquisition scheme with a total of 436 measurements.

Table 4.7: Details of acquisition parameters for different produced acquisition schemes. N describes the number of measurements in the acquisition scheme, δ range/step is the ranges of the pulse width and step size between measurements, and Δ is the minimum and maximum diffusion times. R^2-scores show results from doing ANN fitting of data simulated with these individual acquisition schemes, where R is the radius, (f_{ic}) the intracellular volume fraction, (f_{EES}) the extracellular-extravascular volume fraction and (f_{VASC}) is the vascular volume fraction.

N	**δ range/step [ms]**	**Δ [ms]**	**max b** $[s/mm^2]$	**R^2-scores**			
				Radius (R)	f_{ic}	f_{EES}	f_{VASC}
145	10-40/10	3/10	6717	-0.078	0.553	-0.247	0.656
145	10-60/15	5/12	13454	0.271	0.606	-0.462	0.547
172	10-50/10	8/15	18549	0.328	0.642	0.084	0.783
436	10-60/5	5/12	14773	0.552	0.777	0.512	0.905

Figure 4.7 illustrates the outcomes of training with a significant level of noise. Similarly to training without noise, (a) shows the parameter predictions on the new acquisition scheme with 145 measurements, while 4.7 (b) shows the predictions on the extended acquisition scheme with higher b-values but the same number of measurements. (c) Shows the extended version with both more measurements and higher b-values compared to (a). Furthermore, (d) shows predictions from training with an acquisition scheme with both more measurements and higher b-values. Table 4.8 explains the details of the four different acquisition schemes. Training time for all models was under 700 s.

Table 4.8: Acquisition details and R^2-scores from training on the different extended protocols with added noise. N describes the number of measurements in the acquisition scheme, δ range/step is the ranges of the pulse width and step size between measurements, and Δ is the minimum and maximum diffusion times. R^2-scores show results from doing ANN fitting of data simulated with these individual acquisition schemes, where R is the radius, (f_{ic}) the intracellular volume fraction, (f_{EES}) the extracellular-extravascular volume fraction and (f_{VASC}) is the vascular volume fraction.

N	**δ range/step [ms]**	**Δ [ms]**	**max b** $[s/mm^2]$	**R^2-scores**			
				Radius (R)	f_{ic}	f_{EES}	f_{VASC}
145	10-40/10	3/10	6717	-0.285	0.629	-0.475	0.434
145	10-60/15	5/12	13454	-0.03	0.723	-0.107	0.552
172	10-50/10	8/15	18549	0.102	0.767	0.073	0.602
436	10-60/5	5/12	14773	0.633	0.924	0.830	0.935

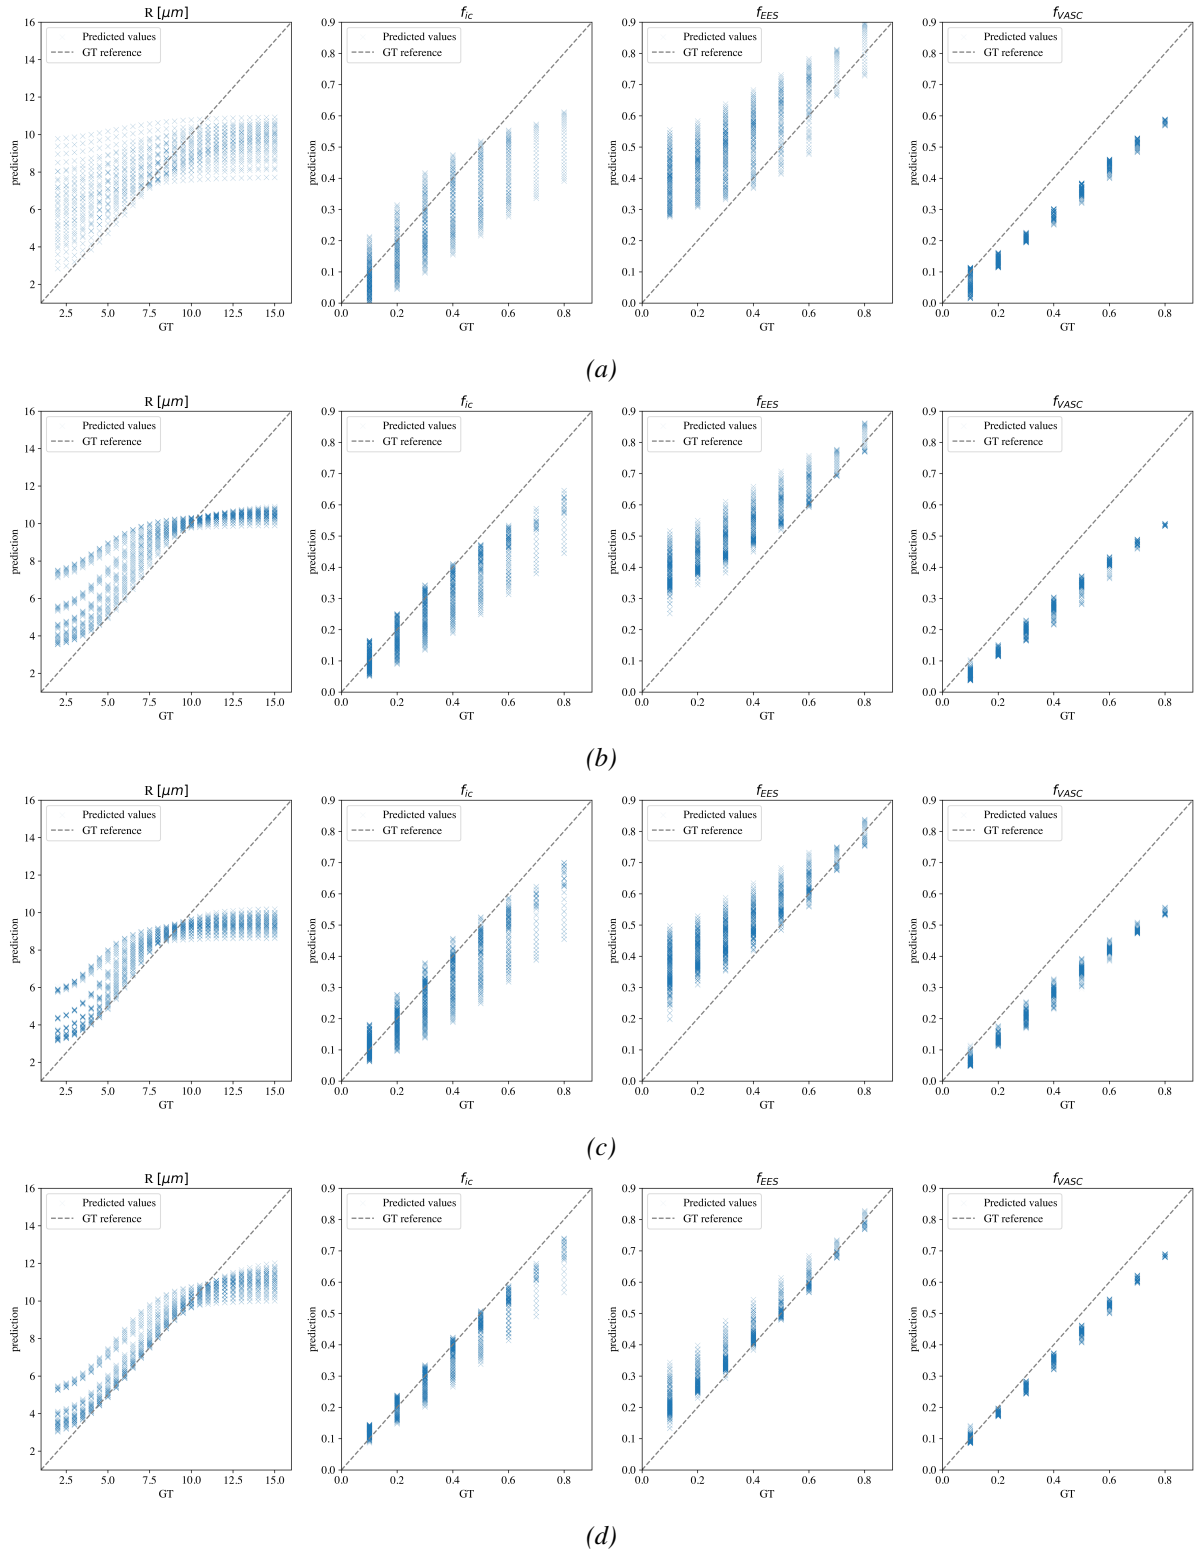

Figure 4.7: Comparison of ground truth (GT) and predicted parameter maps on the training set for the extended protocols with noise. Showing radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}). (a) Predictions from a real acquisition scheme with a total of 145 measurements, (b) scheme extended with larger but less dense δ-range and higher Δ, (c) extension to higher b-values with 172 measurements, (d) a further extension of the acquisition scheme with a total of 436 measurements.

Similarly to what was shown for the fitting of the abbreviated and robust protocol, table 4.9 shows the calculated mean, standard deviation and MSE on four example values (R = 8 $\mu m, f_{ic} = 0.2, f_{EES} = 0.7, f_{VASC} = 0.1$). The MSE is also compared to the calculated MSE on the large-scale model training. This shows that predictions at the lowest have an MSE of 5.9 times the large-scale prediction when predicting intracellular volume fraction with the n = 436 scheme. The highest MSE compared to the large-scale MSE is found when predicting the vascular volume fraction with the n = 172 extended scheme.

Table 4.9: Statistical measures of example values (R = 8 $\mu m, f_{ic} = 0.2, f_{EES} = 0.7, f_{VASC} = 0.1$) on artificial neural network training predictions on the abbreviated and robust dataset. Mean, standard deviation and mean squared error (MSE) on the predicted radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}). Relative MSE error is the calculated MSE compared to MSE on large-scale predictions.

		$R[\mu m]$	f_{IC}	f_{EES}	f_{VASC}
n=145	Mean	8.60 ± 0.86	0.158 ± 0.067	0.752 ± 0.050	0.068 ± 0.030
	MSE	1.088	0.0063	0.0052	0.0019
	MSE/MSE large-scale	54.4	31.3	52.1	322.2
n=145 (ext)	Mean	8.674 ± 0.464	0.145 ± 0.053	0.758 ± 0.031	0.030 ± 0.013
	MSE	0.669	0.0059	0.0043	0.0050
	MSE/MSE large-scale	33.5	29.6	42.8	836.9
n=172	Mean	8.40 ± 0.40	0.152 ± 0.050	0.693 ± 0.024	0.085 ± 0.010
	MSE	0.322	0.0048	0.0006	0.0003
	MSE/MSE large-scale	16.1	24.2	6.2	50.3
n=436	Mean	8.754 ± 0.697	0.186 ± 0.031	0.680 ± 0.016	0.103 ± 0.007
	MSE	1.053	0.0012	0.0006	0.0001
	MSE/MSE large-scale	52.7	5.9	6.3	8.6

4.2 Comparison with Existing methods

In this section, results from comparing the ANN approach is compared with the existing methods in terms of computational time and accuracy.

4.2.1 Comparison of Computational Time

The ANN-based parameter estimation was shown to be more than 150 times faster than the NLLS fitting with DMIPY on the simulated dataset (n = 3466) and close to 300 times faster on the *in vivo* dataset (n = 160) using GPU. The time consumption of all tested methods is shown in table 4.10. For the ANN method, the time used to train the neural network on the simulated signal with the corresponding acquisition scheme is shown together with the time it takes to predict results. From the table, it is clear that the DMIPY method is very slow compared to the two others, but the GPU-accelerated MDT model-fitting method is faster than the ANN.

Table 4.10: Time consumption for the developed ANN, and the two existing methods. Training time is only relevant for the ANN that needs to be retrained for every acquisition protocol.

Framework	Dataset	Device	Training time [s]	Prediction time [s]	Time per voxel [s]
ANN	*in Silico*	GPU	899	< 10	0.097
		CPU	5950	< 10	0.648
	in vivo 160	GPU	693	810	0.001
	in vivo 20	GPU	669	7398	0.010
DMIPY	*in silico*	CPU	0	146 627	15.972
	in vivo 160	CPU	0	410 540	0.390
MDT	*in silico*	GPU	0	38	0.004
	in vivo 160	GPU	0	73	2×10^{-4}
	in vivo 20	GPU	0	259	3×10^{-4}

4.2.2 Accuracy of Existing Methods

Scatter plots showing results from fitting the DMIPY model on large-scale simulated data (n = 3466) are shown in Fig. 4.8. Radius predictions are accurate until reaching 10 μm, from where it does not manage to predict higher values.

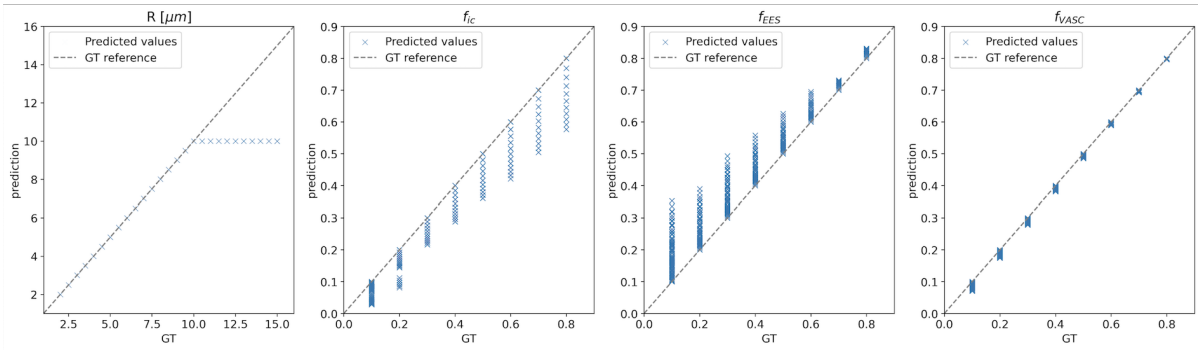

Figure 4.8: Comparison of ground truth (GT) and predicted parameter maps on the simulated large-scale dataset using DMIPY. Showing radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}).

The fitting of the simulated large-scale dataset (n = 3466) with MDT using the Powell optimizer was not successful, as all parameters were predicted to be constant over the entire area. Radius was predicted to be 0.8 μm in all voxels, and the IC, EES and VASC volume fractions were predicted to be -1.117, -1.117 and 3.235, respectively. Results from the second optimizer method, Fig. 4.9, show that it did not manage to make accurate predictions on the *in silico* dataset with n = 3466 measurements.

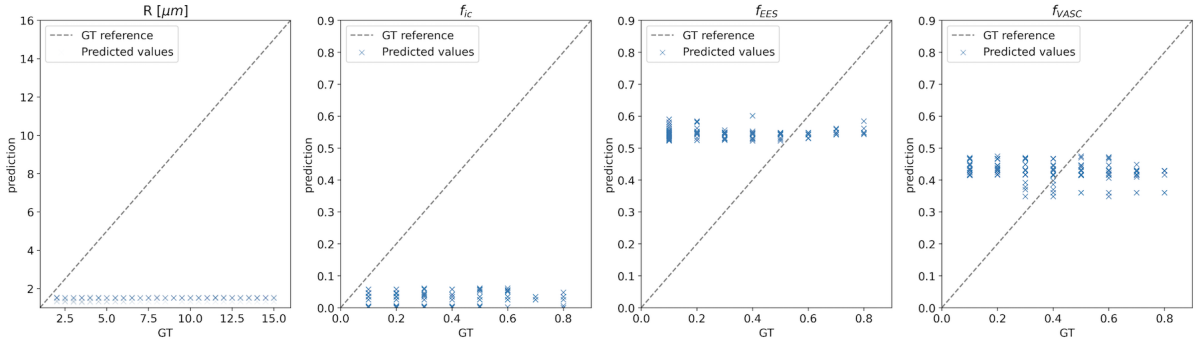

Figure 4.9: Comparison of ground truth (GT) and predictions of parameter maps on the simulated large-scale dataset using MDT. Showing radius (R), intracellular volume fraction (f_{ic}), extracellular-extravascular volume fraction (f_{EES}) and vascular volume fraction (f_{VASC}).

Chapter 5

Discussion and Conclusion

The goal of this master project was to explore whether VERDICT microstructure fitting in glioma tissue could be approached with a self-supervised fully connected neural network. To do this, a neural network was developed and tuned, before it was compared to already existing methods. The network was successfully trained and tested on simulated large-scale data (n = 3466) before being applied to *in vivo* data, returning the four model parameters within expected ranges. However, model fitting on data simulated with *in vivo* acquisition schemes (n = 160, n = 20) did not provide as accurate results as for the large-scale simulated data (n = 3466). For this reason, further datasets were simulated and fitted based on acquisition schemes designed for a high-field MRI protocol, and it was illustrated that these protocols could be better suitable for fitting the VERDICT model. The developed ANN has been compared to existing methods in terms of time consumption and accuracy. The results indicate that the method is more efficient when combining time consumption and accuracy. In this section, the advantages, as well as the limitations of this approach, are discussed.

5.1 Parameter Predictions with Neural Network

5.1.1 Initial Model Fitting on Large-scale Simulated Dataset

This section discusses the results from the initial training process when a large-scale dataset (n = 3466) was simulated and used for model selection. The testing of the model's ability to generalization is also discussed.

Parameter Estimations from Model Training

Figure 4.2 showed that the predictions from the training data did not deviate significantly from the ground truth values. The volume fractions achieved very high accuracy on the training data. Since the ground truth volume fractions only consisted of 10

different values, the trends are slightly harder to determine. However, the first obvious pattern is that the predictions of the vascular component had less deviation from ground truth values and thus more accurate predictions. The deviations were minor but more significant in the mid-range between 0.3 - 0.6. For the IC compartment, predictions were slightly lower than GT values with the most significant deviations in mid-ranges between 0.2 and 0.6. The EES compartment showed the opposite patterns with predictions slightly higher in the same area. The opposite trends were expected as all fractions are normalized to sum to one. The computed example mean values with the corresponding low standard deviation, table 4.3, indicates that the errors are rather systematic and not random, which supports what has been seen visually.

By examining the signal equation for the various compartments, it becomes apparent that the only parameter that is varying in the EES and VASC compartments is the b-value. This implies that the modeling of these compartments disregards the individual acquisition parameters. However, it should be acknowledged that lower b-values result in reduced signals from these compartments, potentially making them more challenging to model.

The predictions of the radius achieved high accuracy, and they were more precise in the mid-ranges, meaning between 5-10 μm. One hypothesis is that the reason for this could be the complexity involved in simulating the spherical compartment, as described in equation 2.19. In the early phases of the study, it was discovered that the model did not manage to compute a good accuracy of simulated data with a high upper range of radius (meaning up to 30 μm which was first used). It was experimented with models trained on data simulated with a maximum of 15 μm, which gave good results. However, it was decided to proceed with a model of up to 20 μm as it was not desirable to exclude some possible expected tissue values. This experiment supports the hypothesis that the computation of the spherical compartment does not manage to compute the high-range radius as well as the lower-range radius.

The formula presented by Balinov et al. [36], i.e. eq. 2.19, might reveal some of the complexity involved in computing the signal from the spherical compartment. As the radius is a squared element in the denominator in the computed sum, high radius values will be equivalent to a smaller sum. However, the radius is also an element in the α element, and the α decreases with an increasing radius. All combined, when the

radius increases the sum becomes bigger, and thus the logarithm of the signal is more negative. This means that the signal from the spherical compartment becomes lower with a high radius value, compared to a higher signal with a lower radius value. This becomes more clear when the radius is exceeding 10 μm. The small amount of signal from the spherical compartment for high radius values might therefore explain the difficulty of predicting the high radius values. Increasing diffusion times and diffusion duration is also components that would result in a lower signal from the spherical compartment, but a higher gradient strength could give a larger signal and thus it may facilitate the computation of the spherical compartment.

Comparing R^2-scores from table 4.2, it is obvious that the model trained on data with added noise achieved a lower R^2-score than the model trained on simulated data without noise. The score of the radius was the one that decreased the most, which also supports that the computation of the spherical compartment is complex and also therefore more sensitive to noise.

The trends seen in predictions of volume fractions can be partly explained by the training data. The GT volume fractions that were used to simulate the data show unbalanced training data, fig. 4.2 (a). This showed that the IC volume fractions contain more data points in lower ranges than in the high ranges, which was also the case for the two remaining volume fractions. This means that the lower ranges included more data points, which in one way would have made the predictions more accurate, but it does also mean that it contained more data points to wrongly predict. The prediction of the IC volume fraction at 0.1 is very accurate, which is the one containing the most data points. The highest values are the ones containing the least amount of data points, and they show less deviation from ground truth in all volume fractions. This is likely to be explained by the low amount of data points and it is thus showing less variation. Overall, this might indicate that training on more data points could have made the model even more accurate.

Because the radius is a parameter in the computation of the spherical component, a logical assumption would be that the volume fraction would compensate for an inaccurate radius estimate by underestimating the IC volume fraction. As the fractions must sum up to one, this would then have affected the other compartments. The computation of the spherical compartment is more complex and involves more components that

could lead to inaccuracies than the other compartments. This supports the assumption that the IC compartment partly causes the inaccuracies of the other volume fractions. Thus, a more efficient way to estimate fractions that need to add up to one simultaneously could be beneficial.

Model Testing

The model testing on a simulated large-scale test set (n = 3466) showed similarly good results as the model tested on the training data. Fig 4.3 (c) showed that when the predicted values were plotted alongside the ground truth reference, the differences between the two were minimal and the predictions followed the same trend as for the parameters obtained during training. The R^2-scores, table 4.4 supported the visually shown results on the test set. In order to determine whether or not the model overfits the training set, additional variability in the data points would be necessary. However, as the accuracy of training data are not higher than that of the test data, this indicates that the model manages to generalize and does not overfit.

5.1.2 Model Training and Prediction for Application on *in vivo* MRI

Training on Abbreviated and Robust Protocol

Going from 3466 to 160 and 20 image volumes, the training loss increased in general, which was expected as the simulated robust and simulated datasets contain fewer data points. The loss of both models converged within the 50 epochs, but the robust dataset had in general a lower training loss than the abbreviated dataset, table 4.5. The table also shows that R^2-scores on the predicted parameter values are significantly lower than on the large-scale dataset (n = 3466), and table 4.6 shows that standard deviations are higher, and MSE is much higher compared to the large-scale fitting.

These trends were expected, as the abbreviated dataset contained one eight the number of data points compared to the robust dataset. Fig 4.4 shows scatter plots of the fit of the predictions for an abbreviated (n = 20) and robust (n = 160) protocol, which is worse than the fits on the large-scale (n = 3466) trained model. Comparing the two results, scatter plots show that predictions are less accurate for the abbreviated protocol. It was expected that the prediction accuracy would decrease with a shallower dataset, as it contained fewer data points per voxel. It is noticeable that the errors of the IC

and EES volume fractions followed the same pattern as the previously discussed result, where the IC compartment was underpredicted and the EES compartment was overpredicted. This supports that these systematic errors are partly caused by the network architecture.

Predictions on Abbreviated and Robust *in vivo* MRI Protocol

From testing the trained models on the robust MRI scan (n = 160), fig. 4.5 (a) and table 4.5, one could see that the predicted radius is overall slightly higher in the glioma tissue, the IC volume fraction is lower and the EES volume fraction is higher. All these trends match the findings in the previous papers described in section 2.2.5, which suggests that the model fitting has produced plausible outcomes. The model trained on the abbreviated protocol (n = 20) did not manage to predict a radius that is compatible with the expected results, which can be assumed to be caused by the lack of information in the short sequence. However, the predictions of the three volume fractions show the same promising patterns as for the fitting on the robust MRI scan (n = 160), with values compatible with the literature [44] [43].

On the robust MRI fitting, a distinct band of high intensity can be seen on the right hemisphere, predicting a smaller radius, higher IC volume fraction and lower EES fraction. From investigating the patient MRI scans, it was seen that the band was not recognizable in the structural scans (i.e. T1 and T2-weighted images), and neither on the diffusion-weighted images with lower b-values, Fig. 4.5. However, as seen in Fig. 3.4 (a), the band is recognizable on the high b-value dMRI. This argues that the band is not a consequence of the model fitting, and rather an artifact from the scanning process. This finding shows that the model fitting is influenced by artifacts on the images, and it has made the model incapable of predicting correct values in the affected area. Fortunately, the voxel-wise nature of the model appears to have prevented artifacts from affecting predictions in other regions of the brain, and the band is recognizable and can be easily classified as incorrect predictions.

5.1.3 Optimization of Acquisition scheme

Results from training the dataset simulated with the rich acquisition scheme showed the potential of making accurate predictions with this method. However, the results from applying this to the acquired *in vivo* sequence show that it does not manage to make accurate predictions on these sequences. To explore this further, Fig. 4.6 shows the

predictions from the model on simulated data on some new acquisition schemes.

Theoretical Approach without Noise

The first fitting 4.6 (a) was done using a real acquisition scheme used on a 9.4 T MR scanner, which could, in theory, cope with larger b-values and more noise. By comparing the scatter plots from Fig. 4.4 and Fig. 4.6, it can be seen that the predictions on this dataset are slightly better than our already trained network on the robust dataset (n = 160). The new fitting shows values closer to the ground truth compared to the previous robust training, which is also supported by comparing R^2-scores from tables 4.5 and 4.7, where the new training has overall higher R^2-scores. One obvious difference between the two situations is that the gradient strengths are remarkably higher in these extended protocols than in the robust protocol. This implies that a higher gradient strength would provide more information in the context of microstructure fitting. The earlier discussion on the potential benefits of higher gradient strengths for computing the spherical compartment is supported by these findings.

From further exploring the scatter plots from Fig. 4.6 and R^2-scores from table 4.7, one can see that the extended version with less dense but higher b-values shows a better fit, especially for the radius. This argues that higher b-values provide more information and a more accurate model. However, the gradient strength is the same in the two situations, which shows that larger Δ and δ values are also crucial. As opposed to stated earlier, this indicates that the fitting is not necessarily dependent on signals acquired with a high gradient strength.

The extended version with 172 measurements and higher b-values, 4.6 (c), shows better radius predictions, with relatively accurate results up to 8 μm. R^2 scores are also higher than the two previous datasets for all parameters. The further extended version, 4.6 (d), with 436 measurements, manages to predict the radius even more precisely, and is starting to approach the result from the training of the rich protocol dataset.

While it may appear that using larger values of δ and Δ results in a more accurate fit, this is not always the case. When considering the extended protocols and the robust protocol, the IC compartments show a worse fit for high-volume fractions. In order to capture the movements of water molecules in a large sphere, it would not intentionally be necessary with a very long diffusion time, as there should be enough molecules to

distinguish the diffusion process. A very long diffusion time would also leave a very low signal that is harder to model. Likewise will long gradient pulses result in a loss of signal.

High B-values and Noise

Even though these results show the theoretical possibility of making accurate predictions with higher b-values, that also means more noise in the data. Comparing Fig. 4.7 to Fig. 4.6, one can see the same patterns over the different schemes. In this fitting, the extended version with 145 measurements is still remarkably better than the original version with the same amount of measurements. One would expect a lower SNR and thus a signal with less information in the extended version, as it has a higher b-value. However, the model seems to be quite robust to noise, as adding noise levels up to 3-4 SNRs did not break the model. From examining dMRI of cancer with high and low b-values, Tang. et al. [77] concluded that with further development in standardization of image acquisition and analysis, dMRI across the entire b-value spectrum is well positioned to become a powerful surrogate in cancer imaging. This statement supports that the use of high b-values could become applicable in clinical statements.

5.2 Comparison with Existing Methods

The fitting with ANN showed promising results in terms of accuracy, and it is also faster than one out of two existing methods. The utilization of deep learning in imaging diffusion microstructures has in some ways demonstrated better performance than conventional methods used previously. Some reasons for this will be further discussed in this section. It should be taken into account that all models have been fitted on the entire brain, whereas VERDICT is designed for tumor tissue. This might have caused some tissue to be more computationally difficult to predict because the model is not designed for such tissue. As this is done for all three models, it is assumed that such effects can be disregarded when comparing the models.

5.2.1 MDT

The MDT approach showed a very fast fitting, using only a few minutes on the largest dataset. However, even though the general statement from the MDT paper claims that the Powell optimizer outperforms other optimizers [15], this was not the case in this experiment as it did not manage to predict probable tissue values. However, the remaining two optimizers were also tested, whereas only one gave predictions. These predicted values seemed randomly distributed and did not seem to be probable tissue values. The MDT was tested with all combinations of datasets and optimizers, whereas none gave better results than shown in figure 4.9. The missing results on the third possible optimizer due to computational restrictions may have caused better outcomes to be overlooked, but based on the general statement from the MDT paper, it is not likely that this optimizer would have performed better. However, it should be taken into account that the lack of accurate predictions might be caused by data or method-specific problems, and a better outcome for this model could have been overlooked. Although the fitting with MDT was not successful, the process still gave useful results in terms of time comparison, as the ANN approach could be compared to another GPU processing method.

5.2.2 DMIPY

The DMIPY approach achieved highly precise radius predictions of up to 10 μm, with volume fraction estimations comparable to those generated by the ANN. As previously explained, the ANN's calculation of the spherical compartment is inspired by DMIPYs way of implementing the formula for computing the signal from the spherical compart-

ment, which was the one presented by Balinov et al. [36], 2.19. A possible consequence of this is that the predictions exhibit similar behavior, leveling off at a radius of 10 μm. By digging deeper into the DMIPY modeling of the spherical compartment, it was seen that a max and minimum limit to the diameter value was set to $(1e^{-2}, 20)$, meaning a radius limit of $(0.5e^{-2}, 10)$. This explains why the DMIPY method does not predict higher values, and might also support the previously suggested reason that there is a limitation in the signal formula. These observations provide further support for the previous assertion that the spherical compartment includes the most complexity in the modeling.

Additionally, it was observed that the DMIPY fitting was much slower on the simulated dataset than on the *in vivo* data. This implies that some calculations are more time-consuming than others, and possibly that computing the parameters from the simulated protocol is not computationally feasible.

5.3 Methodological Considerations

5.3.1 Data Simulation Considerations

To achieve the desired level of precision in the neural network, it is essential that the data quality is high. Initially, this project sought to test the self-supervised neural network method outlined in [49], which stated that the approach would require retraining for each dataset. The aim was to sample accurate parameters by deliberately overfitting the training data. However, the models displayed good generalization capabilities, prompting a change in the project's objective to investigate the feasibility of creating models that could generalize to unseen data. As a consequence, the network was not originally designed with the intention of training it on a large volume of training data.

By exploring previously conducted studies, all studies have used a dramatically higher amount of training data. Lim et al. [49] and Martin et al. [51] simulated and trained their network on 10^5 signal values, while Diao et al. [52] constructed a dataset containing 4×10^6 signals. In comparison to this study, previous research studies have trained their networks on as much as 40 times more data. However, it's important to note that those models were not specifically designed to fit the VERDICT model. Instead, they were designed to fit simpler models that are less computationally expensive. During the work of this project, the impact of more data points was experienced to be crucial. It started out containing 1024 signal values before it was up-scaled in several steps until 9180 data points. Obviously, this also increased the training time of the models, which was why it was not extended even further. However, an even bigger dataset could potentially increase the performance of the network.

Due to the limited number of data points used for training, it is especially crucial to carefully consider the potential data bias. As illustrated in Fig. 4.2, the parameter values used for simulating the data are predominantly concentrated in the lower ranges of volume fractions, compared to the high ranges. Initially, this was not deemed significant since the initial predictions, Fig. 4.2, demonstrated no difference in accuracy between the low and high-range values. However, later predictions conducted on shallower acquisition schemes, as shown in Fig. 4.6 and 4.7, revealed that the accuracies of the lower range values in volume fractions are higher. This may be due to the lack of ground truth values in the high ranges.

Acquisition Scheme Optimization

In the attempt of finding a more suitable protocol scheme, it was demonstrated that the proposed ANN can produce more suitable but still realistic protocols. However, it would be preferable to go through the possible acquisition schemes in a more systematic way, such as comparing which combinations out of the 3466 measurements would make the best. This could be done by testing all possible combinations of 10 measurements and seeing which combinations that gave the best result.

The extended acquisition schemes are constructed for a 9.4 T MR machine, which is not the clinical standard. This was done to investigate whether one could obtain a decent result with shorter acquisition protocols. However, the results showed that it was not necessarily too dependent on the gradient scheme. Therefore it would be interesting to do a fitting to shorter sequences with higher δ and Δ values but with a clinically available gradient strength.

5.3.2 Network Architecture Considerations

As already discussed, there seems to be a clear correlation between the predicted EES volume fraction and IC volume fraction. A possible reason for this can be the order in which they are computed in the neural network and the way that they are normalized. To simplify the neural network, it would have been possible to reduce the number of volume fractions from three to two since they must sum to one. By doing this, the third compartment could have been calculated outside of the hidden layers as one minus the two calculated fractions, eliminating one parameter that the neural network would need to estimate.

Another trend that needs to be considered is the radius prediction on values higher than 10 μm. Scatter plots of predictions on the rich protocol *in silico* data reveal that the accuracy of the fit is deteriorating for values above 10 μm, as illustrated in Fig. 4.2. This trend is also evident in the extended realistic acquisition schemes, as shown in Fig. 4.6. As outlined in section 3.2, the computation of the spherical compartment that considers the radius is carried out in a similar fashion as in the DMIPY framework. It can be observed in Fig. 4.8 that the DMIPY toolbox also fails to predict radius values exceeding 10 μm. This suggests that the challenges associated with fitting radius values may be caused by the previously discussed complexity of the method used to calcu-

late the spherical compartment. Thus, it could be beneficial to optimize the network in terms of the spherical compartment.

Alternative Arichitectures

As stated earlier, this network architecture was adapted from previously conducted studies [49], [50]. These two studies both consisted of an input layer, three fully connected hidden layers and an output layer, using the ELU activation function and Adam optimizer. However, other similar approaches have taken other choices in terms of architecture.

In an attempt of fitting a diffusion-relaxation model of white matter, Martins. et al. [51] trained various network designs with different numbers of hidden noise. From comparing a network consisting of three hidden layers with a decreasing amount of nodes with a network with four consistent hidden layers, they found that increasing the width or depth of the fully connected ANN was not a promising avenue. This project also explored the use of a greater number of hidden layers, with similar results obtained. This suggests that the proposed layer structure represents the optimal setting and that there is limited potential for improving the selected layer structure. Another promising study conducted by Diao et al. in November 2022 [52] also used an FCNN approach for comparison purposes, consisting of a decreasing amount of nodes in the hidden layers. This suggests that it could be interesting to test a decreasing amount of nodes in the suggested ANN, but significant improvement may not be expected based on the findings of Martins et al.'s study.

The main purpose of Diao et al.'s study was to propose a new method based on the encoder-decoder recurrent neural network (RNN) in an attempt to replace the conventional NLLS fitting and to challenge the FCNN approach. In May 2023, Zheng et al. proposed a framework based on a transformer network structure and introduced an inductive bias-model bias into the transformer with the use of a sparse coding technique to facilitate the training process. Both these two proposed methods showed promising results in terms of fitting their respective microstructure models. These completely different network architectures, therefore, appear interesting to investigate for fitting the VERDICT model in gliomas. On the other hand, Valindra et al. [78] explored three

different deep regression networks for fitting VERDICT including one FCNN and two deeper networks consisting of more complex hidden layers. The findings argued that the simple FCNN approach is adequate to estimate parametric maps on models like the classic VERDICT.

Martins et al. [51] also applied the Adam optimizer and MSE as the loss function, and as far as can be determined no other optimizer or loss function has been applied successfully in other studies. Martins et al.'s study, however, employed the hyperbolic tangent (tanh) activation function, in contrast to the ELU used in this project. Diao et al. [52] applied the LeakyRELU activation function. These are optimizers that could also be tested in terms of fitting the VERDICT model. Nonetheless, significant efforts were made in the model selection phase, making it unlikely that remarkable improvements in the results would be achieved by implementing these various optimizers.

5.3.3 Choice of Microstructure Model

The choice of microstructure model was done based on the result from Zaccagna et al. [43], who demonstrated a good fit of the VERDICT model in glioma tissue compared to histological results. However, one can not rule out that some of the issues encountered in this project might also be due to the design of the VERDICT model. One can argue that the three compartments are not adequate in predicting all seven tissue parameters. However, further analysis is necessary for investigating this.

Furthermore, it has been found that tumor cells can infiltrate centimeters away from the contrast-enhancing mass on the MRI scans [79] [80]. In such a case, it is natural to consider the potential inadequacy of the VERDICT modeling technique in capturing the intricate composition of such a diverse and intricate tissue. Additionally, the original VERDICT model was designed to compute the IC, EES and vascular spaces in an animal model of prostate cancer at a field strength of 9.4 T which has higher diffusion gradient capabilities [43]. This might partly explain the difficulties concerning the computational modeling including noise at clinical field strengths, meaning data acquired with a field strength of 3 T.

5.4 Ethical Considerations

To ensure the integration of this self-supervised deep learning approach as a supplementary or alternative method to conventional biopsy, it is important to address the ethical considerations. One considerations concern issues related to epistemic and misguided, inconclusive evidence [69]. One epistemic concern regarding this project is the accuracy of the predictions. It has been shown that the accuracy of predictions is highly dependent on the data, which means that using this approach on unsuitable data could potentially lead to misguided, inconclusive or inscrutable evidence. Even in scenarios where the approach demonstrates near-perfect performance, it is crucial to investigate the precision of the predictions. Ideally, thorough comparisons to histological estimations of tissue parameters are necessary to ensure that this approach could perform just as well as the traditional method. However, comparison to gold standards in this field is methodologically challenging. Most importantly, the method should be applied only in circumstances where it is well-tested and proven to be appropriate.

Another consideration to address is related to normative and unfair outcomes [69]. If the model is trained on biased data, by for example using data with too low radius values, one could risk that the model expects normative behavior and does not capture extreme or severe cases. It has been addressed that the developed model does not manage to predict very high radius values higher than 15 μm, and even though such values are not expected, it shows that the model does to some extent address the issue of normative behavior. This could be severe in the cases of very large glioma cells. For this reason, consideration of data in the training process becomes extra important.

Even though the predictions of the network might be reliable, the issue of traceability still arises [69]. It is important that the algorithm is implemented in such a way that it provides traceability and transparency. However, the question about how the neural network does its predictions is still present. Sufficient knowledge about the network structure and machine learning technicalities is needed. It must still be considered that the traceability might be limited compared to the conventional method. Another concern regarding traceability is the algorithm's ability to function properly even if the measured diffusion data changes over time, for example, due to scanner upgrades. Continuous evaluation of the algorithm is likely necessary.

This tool is designed to complement human evaluations and should be utilized by individuals as a support, rather than replacing their critical judgments. Nevertheless, these discussed issues are important to take into consideration when implementing such an algorithm. With further development of this AI implementation into diffusion microstructure imaging, it becomes essential to address these concerns related to epistemic considerations, normative behavior, and traceability.

5.5 Conclusions and Future Work

This project has successfully demonstrated that the self-supervised ANN approach is suitable for complex VERDICT microstructure fitting in glioma tissue. While it has previously been demonstrated for simpler methods, the complex case of VERDICT fitting in gliomas had not been demonstrated until now. Consequently, this research has the potential to significantly improve the accessibility of glioma treatment response assessment. Through extensive training on large-scale simulated data, the network exhibits the potential of a high level of confidence in the complex predictions of the cell radius, as well as the IC, EES, and vascular volume fractions. Notably, the method's applicability to *in vivo* dMRI data has been successfully demonstrated. As compared to existing methods, the proposed ANN showed a better performance when taking both time consumption and accuracy into consideration. It is worth emphasizing that further optimization of the ANN approach was not pursued within the scope of this project, considering its challenging nature.

As expected, the analysis of the relationship between model fitting and acquisition protocol revealed that more detailed sequences were more suitable for the model fitting, while shorter and less detailed protocols resulted in inaccurate predictions. This method is not recommended to be used in image protocols with less than 20 measurements due to the lack of sufficient information for successful fitting. For predicting a radius of up to 10 μm, 145 high b-value measurements could be sufficient, but n = 436 measurements provide more reliable predictions in higher radius ranges as well. From experimenting with higher b-values and larger ranges of diffusion, it can be concluded that the model fitting gives better results. With the support of Tang et. al. [77], the use of the entire b-value spectrum in dMRI cancer analysis is a promising field.

Future work could include optimizing the self-supervised network in terms of even more complex signal simulation such as mixed signals from the spherical compartments. Also training the network on larger amounts of data with a more balanced dataset would be important. Furthermore, it would be interesting to do a more detailed investigation of gradient combinations and acquisition schemes. Finally, it will be important to validate this VERDICT fitting approach by investigating a glioma tissue structure with both this proposed ANN method and histological pathologies.

Bibliography

[1] H. Hyare, S. Thust, and J. Rees, "Advanced mri techniques in the monitoring of treatment of gliomas," *Current Treatment Options in Neurology*, vol. 19, pp. 1–15, 3 2017. 1

[2] D. N. Louis, A. Perry, G. Reifenberger, A. von Deimling, D. Figarella-Branger, W. K. Cavenee, H. Ohgaki, O. D. Wiestler, P. Kleihues, and D. W. Ellison, "The 2016 world health organization classification of tumors of the central nervous system: a summary," *Acta Neuropathologica*, vol. 131, no. 6, pp. 803–820, Jun 2016. 1

[3] A. Daher and J. de Groot, "Rapid identification and validation of novel targeted approaches for glioblastoma: A combined ex vivo-in vivo pharmaco-omic model," *Experimental Neurology*, vol. 299, pp. 281–288, 2018, experimental Approaches to Advance Therapies for Tumors of the Central and Peripheral Nervous System. 1

[4] K. Anjum, B. I. Shagufta, S. Q. Abbas, S. Patel, I. Khan, S. A. A. Shah, N. Akhter, and S. S. ul Hassan, "Current status and future therapeutic perspectives of glioblastoma multiforme (gbm) therapy: A review," *Biomedicine and Pharmacotherapy*, vol. 92, pp. 681–689, 2017. 1

[5] K.-J. Langen, N. Galldiks, E. Hattingen, and N. J. Shah, "Advances in neuro-oncology imaging," *Nature Reviews Neurology*, vol. 13, no. 5, pp. 279–289, May 2017. 1

[6] O. M. Henriksen, M. D. Mar, Álvarez Torres, P. Figueiredo, G. Hangel, V. C. Keil, R. E. Nechifor, F. Riemer, K. M. Schmainda, E. A. H. Warnert, E. C. Wiegers, T. C. Booth, N. Sollmann, H. Poptani, and P. W. Hales, "High-grade glioma treatment response monitoring biomarkers: A position statement on the evidence supporting the use of advanced mri techniques in the clinic, and the lat-

est bench-to-bedside developments. part 1: Perfusion and diffusion techniques," *Article 810263 1 Front. Oncol*, vol. 12, p. 810263, 2022. 1

[7] D. Le Bihan, E. Breton, D. Lallemand, M. L. Aubin, J. Vignaud, and M. Laval-Jeantet, "Separation of diffusion and perfusion in intravoxel incoherent motion mr imaging." *Radiology*, vol. 168, no. 2, pp. 497–505, 1988, pMID: 3393671. 1, 2.1.2, 2.2.2

[8] D. C. Alexander, P. L. Hubbard, M. G. Hall, E. A. Moore, M. Ptito, G. J. Parker, and T. B. Dyrby, "Orientationally invariant indices of axon diameter and density from diffusion mri," *NeuroImage*, vol. 52, no. 4, pp. 1374–1389, 2010. 1

[9] T. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith, "Characterization and propagation of uncertainty in diffusion-weighted mr imaging," *Magnetic Resonance in Medicine*, vol. 50, no. 5, pp. 1077–1088, 2003. 1, 2.2.2, 2.2.3

[10] S. D. Vasylechko, S. K. Warfield, O. Afacan, and S. Kurugol, "Self-supervised ivim dwi parameter estimation with a physics based forward model," *Magnetic Resonance in Medicine*, vol. 87, no. 2, pp. 904–914, 2021. 1, 2.3.1

[11] 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*, vol. 61, no. 4, pp. 1000–1016, 2012. 1

[12] E. Panagiotaki, S. Walker-Samuel, B. Siow, S. P. Johnson, V. Rajkumar, R. B. Pedley, M. F. Lythgoe, and D. C. Alexander, "Noninvasive Quantification of Solid Tumor Microstructure Using VERDICT MRI," *Cancer Research*, vol. 74, no. 7, pp. 1902–1912, 03 2014. 1, 1.1, 2.2.3, 2.2.3, 2.2.3, 2.2.3, 3.1.3, 3.4

[13] E. Panagiotaki, R. W. Chan, N. Dikaios, H. U. Ahmed, J. O'Callaghan, A. Freeman, D. Atkinson, S. Punwani, D. J. Hawkes, and D. C. Alexander, "Microstructural characterization of normal and malignant human prostate tissue with vascular, extracellular, and restricted diffusion for cytometry in tumours magnetic resonance imaging," *Investigative Radiology*, vol. 50, no. 4, pp. 218–227, 2015. 1.1, 2.2.3, 2.4, 2.2.5

[14] E. Panagiotaki, T. Schneider, B. Siow, M. Hall, M. Lythgoe, and D. Alexander, "Compartment models of the diffusion mr signal in brain white matter: a taxon-

omy and comparison," *Neuroimage*, vol. 59, no. 3, pp. 2251–54, 2012. 1.1, 2.2.2, 2.2.3

[15] R. L. Harms and A. Roebroeck, "The maastricht diffusion toolbox (mdt): Modular, gpu accelerated, dmri microstructure modeling," *ISMRM 2017*, 2017. 1.1, 2.2.4, 3.2.4, 3.3.1, 5.2.1

[16] E. Bonet-Carne, E. Johnston, A. Daducci, J. G. Jacobs, A. Freeman, D. Atkinson, D. J. Hawkes, S. Punwani, D. C. Alexander, and E. Panagiotaki, "Verdict-amico: Ultrafast fitting algorithm for non-invasive prostate microstructure characterization," *NMR in Biomedicine*, vol. 32, no. 1, p. e4019, 2019, e4019 NBM-18-0071.R2. 1.1

[17] F. Grussu, M. Battiston, M. Palombo, T. Schneider, C. A. M. G. Wheeler-Kingshott, and D. C. Alexander, "Deep learning model fitting for diffusion-relaxometry: A comparative study," *bioRxiv*, pp. 159–172, 2021. 1.1

[18] M. Cercignani, N. Dowell, and P. Tofts, Eds., *Quantitative MRI of the Brain*. CRC Press, January 2018. 2.1

[19] R. W. Brown, Y.-C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan, *Magnetic Resonance Imaging*. John Wiley & Sons, Ltd, 2014, ch. 1, pp. 1–17. 2.1

[20] E. R. Grüner, "Compendium phys212: Medical physics and technology," *Department of Physics and Technology, University of Bergen*, vol. 1.0, 2012. 2.1

[21] R. M. Henkelman, "Measurement of signal intensities in the presence of noise in mr images," *Medical Physics*, vol. 12, no. 2, pp. 232–233, 1985. 2.1.1

[22] H. Gudbjartsson and S. Patz, "The rician distribution of noisy mri data," *Magnetic Resonance in Medicine*, vol. 34, no. 6, pp. 910–914, 1995. 2.1.1

[23] S. O. Rice, "Mathematical analysis of random noise," *The Bell System Technical Journal*, vol. 23, no. 3, pp. 282–332, 1944. 2.1.1

[24] D. K. Jones, *Introduction and concepts*. Oxford University Press, 2011, p. 5762. 2.1.2, 2.1.2

[25] L. G. Hanson, "Flow and diffusion weighing," *DTU Health Tech, Technical University of Denmark*, vol. 1.0, 2022. 2.1.2

[26] E. O. Stejskal and J. E. Tanner, "Spin diffusion measurements: Spin echoes in the presence of a timedependent field gradient," *The Journal of Chemical Physics*, vol. 42, no. 1, pp. 288–292, 1965. 2.1.2, 2.1.2

[27] H. C. Torrey, "Bloch equations with diffusion terms," *Phys. Rev.*, vol. 104, pp. 563–565, Nov 1956. 2.1.2, 2.1.2

[28] D. Le Bihan and E. Breton, "In vivo magnetic resonance imaging of diffusion," *Comptes Rendus des Seances de l'Academie des Sciences*, vol. 301, no. 15, pp. 1109–1112, Nov 1985. 2.1.2

[29] D. Le Bihan, "Molecular diffusion nuclear magnetic resonance imaging," *Magnetic resonance quarterly*, vol. 7, no. 1, p. 130, January 1991. 2.1.2

[30] H. Seo, K.-H. Chang, D. Na, B. Kwon, and D. Lee, "High b-value diffusion (b = 3000 s/mm2) mr imaging in cerebral gliomas at 3t: Visual and quantitative comparisons with b = 1000 s/mm2," *American Journal of Neuroradiology*, vol. 29, no. 3, pp. 458–463, 2008. 2.1.2

[31] M. Yoshikawa, S. Ohsumi, and S. e. a. Sugata, "Relation between cancer cellularity and apparent diffusion coefficient values using diffusion-weighted magnetic resonance imaging in breast cancer," *Radiat Med*, vol. 26, pp. 222–226, 2008. 2.2

[32] Z. Rumboldt, D. Camacho, D. L. C. Welsh, and M. Castillo, "Apparent diffusion coefficients for differentiation of cerebellar tumors in children," *American Journal of Neuroradiology*, 2006. 2.2

[33] P. Basser, J. Mattiello, and D. LeBihan, "Mr diffusion tensor spectroscopy and imaging," *Biophysical Journal*, vol. 66, no. 1, pp. 259–267, 1994. 2.2.1

[34] D. Le Bihan, J.-F. Mangin, C. Poupon, C. A. Clark, S. Pappata, N. Molko, and H. Chabriat, "Diffusion tensor imaging: Concepts and applications," *Journal of Magnetic Resonance Imaging*, vol. 13, no. 4, pp. 534–546, 2001. 2.2.1

[35] U. Ferizi, T. Schneider, and E. Panagiotaki, "A ranking of diffusion mri compartment models with in vivo human brain data," *Magn Reson Med.*, vol. 72, no. 6, pp. 1785–1792, 2014. 2.2.2, 2.2.3

[36] B. B., J. B., L. P., and S. O., "The nmr self-diffusion method applied to restricted diffusion. simulation of echo attenuation from molecules in spheres and between

planes," *Journal of Magnetic Resonance, Series A*, vol. 104, no. 1, pp. 17–25, 1993. 2.2.3, 5.1.1, 5.2.2

[37] C. H. Neuman, "Spin echo of spins diffusing in a bounded medium," *The Journal of Chemical Physics*, vol. 60, no. 11, pp. 4508–4511, 1974. 2.2.3

[38] J. S. Murday and R. M. Cotts, "Selfdiffusion coefficient of liquid lithium," *The Journal of Chemical Physics*, vol. 48, no. 11, pp. 4938–4945, 1968. 2.2.3

[39] R. H. J. Fick, D. Wassermann, and R. Deriche, "The dmipy toolbox: Diffusion mri multi-compartment modeling and microstructure recovery made easy," *Frontiers in Neuroinformatics*, vol. 13, 2019. 2.2.3, 2.2.4

[40] H. Farooq, J. Xu, and J. e. a. Nam, "Microstructure imaging of crossing (mix) white matter fibers from diffusion mri." *Sci Rep*, vol. 6, 2016. 2.2.4

[41] M. J. D. Powell, "An efficient method for finding the minimum of a function of several variables without calculating derivatives," *The Computer Journal*, vol. 7, no. 2, pp. 155–162, 01 1964. 2.2.4

[42] R. Harms, F. Fritz, A. Tobisch, R. Goebel, and A. Roebroeck, "Robust and fast nonlinear optimization of diffusion mri microstructure models," *NeuroImage*, vol. 155, pp. 82–96, 2017. 2.2.4

[43] F. Zaccagna, F. Riemer, A. N. Priest, M. A. McLean, K. Allinson, J. T. Grist, C. Dragos, T. Matys, J. H. Gillard, C. Watts, S. J. Price, M. J. Graves, and F. A. Gallagher, "Non-invasive assessment of glioma microstructure using verdict mri: correlation with histology," *European Radiology*, vol. 29, pp. 5559–5566, 10 2019. 2.2.5, 2.5, 2.2.5, 3.1.2, 3.4, 5.1.2, 5.3.3

[44] T. A. Roberts, H. Hyare, B. Hipwell, A. Ianus, J. O. Breen-Norris, E. Panagiotaki, D. Atkinson, S. Punwani, J. Rees, S. Brandner, D. C. Alexander, and S. Walker-Samuel, "Quantifying microstructure in low and high-grade brain tumours using verdict mri," *ISMRM 2017*, 2017. 2.2.5, 2.2.5, 5.1.2

[45] T. M. Mitchell, *Machine Learning*. McGraw Hill, 1997. 2.3

[46] M. L. Giger, "Machine learning in medical imaging," *Journal of the American College of Radiology*, vol. 15, no. 3, Part B, pp. 512–520, 2018, data Science: Big Data Machine Learning and Artificial Intelligence. 2.3

[47] B. J. Erickson, P. Korfiatis, Z. Akkus, and T. L. Kline, "Machine learning for medical imaging," *RadioGraphics*, vol. 37, no. 2, pp. 505–515, 2017, pMID: 28212054. 2.3.1

[48] A. Zhang, Z. C. Lipton, M. Li, and A. J. Smola, "Dive into deep learning," *arXiv preprint arXiv:2106.11342*, 2021. 2.3.1, 2.3.2, 2.3.2

[49] J. P. Lim, S. B. Blumberg, N. Narayan, S. C. Epstein, D. C. Alexander, M. Palombo, and P. J. Slator, "Fitting a directional microstructure model to diffusion-relaxation mri data with self-supervised machine learning," in *Computational Diffusion MRI: 13th International Workshop, CDMRI 2022, Held in Conjunction with MICCAI 2022, Singapore, Singapore, September 22, 2022, Proceedings*. Springer, 2022, pp. 77–88. 2.3.1, 2.3.2, 3.2, 3.2.1, 3.2.2, 5.3.1, 5.3.2

[50] S. Barbieri, O. J. Gurney-Champion, R. Klaassen, and H. C. Thoeny, "Deep learning how to fit an intravoxel incoherent motion model to diffusion-weighted mri," *Magnetic Resonance in Medicine*, vol. 83, no. 1, pp. 312–321, 2020. 2.3.1, 2.3.2, 5.3.2

[51] J. P. de Almeida Martins, M. Nilsson, B. Lampinen, M. Palombo, P. T. While, C.-F. Westin, and F. Szczepankiewicz, "Neural networks for parameter estimation in microstructural mri: Application to a diffusion-relaxation model of white matter," *NeuroImage*, vol. 244, p. 118601, 2021. 2.3.1, 2.3.2, 5.3.1, 5.3.2

[52] Y. Diao and I. Jelescu, "Parameter estimation for wmti-watson model of white matter using encoderdecoder recurrent neural network," *Magnetic Resonance in Medicine*, vol. 89, no. 3, pp. 1193–1206, 2023. 2.3.1, 5.3.1, 5.3.2

[53] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, 2015. 2.3.2, 2.3.2

[54] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, p. 8490, may 2017. 2.3.2

[55] C. Farabet, C. Couprie, L. Najman, and Y. LeCun, "Learning hierarchical features for scene labeling," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1915–1929, 2013. 2.3.2

[56] J. J. Tompson, A. Jain, Y. LeCun, and C. Bregler, *Joint Training of a Convolutional Network and a Graphical Model for Human Pose Estimation*. Curran Associates, Inc., 2014, vol. 27. 2.3.2

[57] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, *Going Deeper With Convolutions*. CVPR, June 2015. 2.3.2

[58] M. Helmstaedter, K. Briggman, and S. e. a. Turaga, "Connectomic reconstruction of the inner plexiform layer in the mouse retina," *Nature*, vol. 500, pp. 168–174, 2013. 2.3.2

[59] F. Chollet, *Dive Into Deep Learning*. Simon and Schuster, 2021, ch. The "deep" in deep learning. 2.3.2, 2.3.2

[60] "Differentiation: Composite, implicit, and inverse functions." [Online]. Available: https://www.khanacademy.org/math/ap-calculus-ab/ab-differentiation-2-new 2.3.2

[61] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proceedings of COMPSTAT'2010*, Y. Lechevallier and G. Saporta, Eds. Heidelberg: Physica-Verlag HD, 2010, pp. 177–186. 2.3.2

[62] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017. 2.3.2

[63] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research 1*, vol. 521, pp. 2121–2159, 2011. 2.3.2

[64] N. Srivastava, C. Hinton, A. Krizhevsky, I. Sutskevera, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of machine learning research : JMLR.*, vol. 15, no. 1, p. 1929, 2014. 2.3.2

[65] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ser. ICML'10. Madison, WI, USA: Omnipress, 2010, p. 807814. 2.3.2

[66] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and accurate deep network learning by exponential linear units (elus)," *arXiv preprint arXiv:1511.07289*, 2015. 2.3.2

[67] M. P. T. Kaandorp, S. Barbieri, R. Klaassen, H. W. M. van Laarhoven, H. Crezee, P. T. While, A. J. Nederveen, and O. J. Gurney-Champion, "Improved unsupervised physics-informed deep learning for intravoxel incoherent motion modeling and evaluation in pancreatic cancer patients," *Magnetic Resonance in Medicine*, vol. 86, no. 4, pp. 2250–2265, 2021. 2.3.2

[68] E. Zuccarelli, "Performance metrics in machine learning part 2: Regression," 2021. 2.3.3

[69] J. Morley, C. C. Machado, C. Burr, J. Cowls, I. Joshi, M. Taddeo, and L. Floridi, "The ethics of ai in health care: A mapping review," *Social Science & Medicine*, vol. 260, p. 113172, 2020. 2.3.4, 5.4

[70] J. L. Andersson and S. N. Sotiropoulos, "An integrated approach to correction for off-resonance effects and subject movement in diffusion mr imaging," *NeuroImage*, vol. 125, pp. 1063–1078, 2016. 3.1.2

[71] R. Fick, R. Deriche, and D. Wassermann, "The dmipy toolbox: Diffusion mri multi-compartment modeling and microstructure recovery made easy," 2019. 3.2.1, 3.3.1

[72] G. Van Rossum and F. L. Drake, *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009. 3.2.1

[73] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035. 3.2.1, 3.2.2, 3.2.4

[74] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011. 3.2.2

[75] E. Garyfallidis, M. Brett, B. Amirbekian, A. Rokem, S. Van Der Walt, M. Descoteaux, and I. Nimmo-Smith, "Dipy, a library for the analysis of diffusion mri data," *Frontiers in Neuroinformatics*, vol. 8, 2014. 3.3.1

[76] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020. 3.3.1

[77] L. Tang and X. J. Zhou, "Diffusion MRI of cancer: From low to high b-values," *Journal of Magnetic Resonance Imaging*, vol. 49, no. 1, pp. 23–40, Oct. 2018. 5.1.3, 5.5

[78] V. Valindria, M. Palombo, E. Chiou, S. Singh, S. Punwani, and E. Panagiotaki, "Synthetic q-space learning with deep regression networks for prostate cancer characterisation with verdict," in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2021, pp. 50–54. 5.3.2

[79] Q. Zeng, C. Ling, F. Shi, F. Dong, B. Jiang, and J. Zhang, "Glioma infiltration sign on high b-value diffusion-weighted imaging in gliomas and its prognostic value," *Journal of Magnetic Resonance Imaging*, vol. 48, no. 3, pp. 643–651, 2018. 5.3.3

[80] P. J. Kelly, C. Daumas-Duport, D. B. Kispert, B. A. Kall, B. W. Scheithauer, and J. J. Illig, "Imaging-based stereotaxic serial biopsies in untreated intracranial glial neoplasms," *Journal of Neurosurgery*, vol. 66, no. 6, pp. 865 – 874, 1987. 5.3.3

Appendix A

The full source code for this thesis is publicly available on GitHub: https://github.com/theavage/Master-project

Within the repository, you can find the following folder structure with its content:

- **Network** - Contains all code used for running the neural network. All scripts can be run through the main.py file, where it can be specified if the goal is to train the neural network or to predict values from a trained model.
  - Results: Prediction results from a range of training and prediction processes.
  - Models: Trained models
  - dataset.py - For loading dataset
  - main.py - Runs all codes
  - model.py - Defines model architecture
  - train.py - Runs the training loop
  - predict.py - Runs the prediction protocol
  - utils.py - Helper functions
- **DMIPY** - All code used to implement DMIPY method
  - results - Predictions from DMIPY fitting
  - DMIPY_VERDICT.ipynb - Fits VERDICT with DMIPY model
- **Utils** - Scripts for creating scheme files, simulating data, adding noise to data, and visualizing.
  - Bvecs_and_bvals: script for extracting b-value vectors for use in the MDT modeling.
  - b0_norm.py - Script for normalizing images to b0 image.

  - make_schemefile.py - generates scheme file
  - noise.ipynb - Adds noise to simulated signals
  - visualization.ipynb - Visualizes results

- Data - Contains simulated schemes and data