



LUND
UNIVERSITY

Master of Science Thesis
VT2024

High-Resolution Diffusion-MRI by Super-Resolution Reconstruction Using Slice Excitation with Random Overlap

Felix Mortensen

Supervisors

Filip Szczepankiewicz, Malwina Molendowska,
Jakub Jurek and Frederik Testud

Medical Radiation Physics, Lund
Medical Physics Programme
Faculty of Science
Lund University, Sweden
www.msf.lu.se

Populärvetenskaplig sammanfattning

Magnetresonanstomografi, eller MRI, är en central och växande del av dagens vård, vilket ger läkare en bred verktygslåda för att diagnostisera en mängd olika medicinska tillstånd. En typ av undersökning som går att utföras med en magnetkamera är diffusionsundersökningar, som speglar partiklars slumpmässiga rörelse i vatten. Denna typ av undersökning möjliggör kartläggning av organ baserat på vävnadernas mikrostrukturer och egenskaper. Sådan typ av kartläggning öppnar möjligheter för att exempelvis detektera cancer.

I alla MRI-bilder finns brus, och det finns alltid en minimnivå av bruset som kallas för brusgolvet. Diffusionsundersökningar är inte en lätt undersökning att utföra rent tekniskt, och den insamlade signalen är ofta starkt påverkad av detta brusgolv. Det är därför viktigt att man samlar in så mycket signal som möjligt, för att minimera effekterna som brusgolvet har på signalen. Ett sätt att göra detta är att samla in signalen med tjockare voxlar. Detta gör att signalen i varje voxel ökar, och brusgolvseffekten minskar. Nackdelen med detta är dock att upplösningen i den riktning som man utökat voxelstorleken minskar. Ett annat sätt öka signalen i förhållande till bruset är utföra experimentet flera gånger, och ta ett medelvärde av de insamlade signalerna. Nackdelen med detta är att undersökningstiden ökar med varje repetition. En MRI-undersökning är alltså alltid en vägning mellan signal-till-brus förhållandet, upplösning och insamlingstid. Superupplösning har under de senare åren visat sig vara en lovande lösning till detta problem, där man kan bibehålla en hög upplösning och förbättrat signal-till-brus förhållande utan att dramatiskt påverka undersökningstiden.

I denna studie testas en ny superupplösningsmetod, SERO, och jämförs med redan etablerade metoder (den så kallade Slice Shifting-metoden och konventionell högupplöst signalinsamling). Metoden går ut på att samla in ett stort antal lågupplösta MRI-bilder i en slumpmässig ordning, som sedan återskapas till fyra högupplösta parameterkartor som representerar vävnadens egenskaper.

Simuleringar gjordes först i en dimension följt av rekonstruktion på ett numeriskt hjärnfantom. Metodens precision och träffsäkerhet evaluerades och jämfördes med redan etablerade metoder. En pulssekvens kompatibel med vår superupplösningsmetod designades och testades på en frisk hjärna, och jämfördes med de tidigare etablerade metoderna i form av hur väl parameterkartorna återskapades. Slutligen designades ett artificiellt nätverk för att skatta parameterkartorna, som sedan jämfördes med minsta-kvadrat anpassningen.

Resultaten från simuleringarna tydde på att vår superupplösningsmetod hade en förbättrad träffsäkerhet än de metoder som den jämfördes med. Precisionen hos den konventionella högupplösta signalinsamlingsmetoden var dock bättre än vår superupplösningsmetod. Resultatet från det artificiella nätverket var positivt i en dimension, där nätverket både hade en förbättrad träffsäkerhet och precision i jämförelse med en minsta-kvadrat anpassning. Däremot var situationen omvänd när nätverket applicerades på riktig data från en frisk hjärna, där nätverket inte lyckades återskapa några korrekta parameterkartor.

Abstract

Magnetic Resonance Imaging (MRI) is a powerful tool for diagnostics in medical imaging, where diffusion-weighted (DW) MRI is a modality used for probing tissue microstructure and function at a molecular level.

Due to poor signal-to-noise ratio (SNR) and spatial resolution in DW-MRI, the noise floor typically induces a large signal bias by elevating the signal intensity. This bias affects the accuracy of the experimental result. A certain level of SNR is needed to minimise the noise floor effect. The purpose of this study was to solve issues regarding the noise floor signal bias problem through a novel sampling scheme and super-resolution reconstruction (SRR).

In recent years, SRR has been a promising solution to the problem. SRR ensures that the resolution is recovered in the through-plane direction while maintaining a high SNR, without affecting the acquisition time much. This is done by acquiring a large quantity of low-resolution images of the object where each image contributes with new information, and finding a single high-resolution image that explains these low-resolution images. This is however typically an ill-posed problem in a least squares sense.

In this study, a novel approach to SRR is proposed with Slice Excitation with Random Overlap (SERO) instead of serial slice excitation, and compared to the established Slice Shifting method and conventional high-resolution direct sampling. Sampling data by SERO results in a distribution of repetition times (TR), ensuring variable intra-slice T_1 -weighting, helping the ill-posedness of the SRR-problem.

A reconstruction algorithm utilising numerical gradient decent to estimate diffusion and relaxation parameters for SERO was implemented. The accuracy and precision of SERO was then assessed through numerical simulations by generating a one-dimensional structure, and through reconstruction of a numerical brain phantom. The accuracy and precision was compared to that of Slice Shifting and high-resolution direct sampling. A machine learning approach for parameter estimation was designed, aimed at increasing the computational time of the fitting process.

An MRI pulse sequence compatible with random slice positions was designed using the Matlab-based Pulseseq framework in order to demonstrate SERO on in vivo brain data.

The results showed an improved accuracy for SERO compared to Slice Shifting and high-resolution direct sampling in the lower SNR-regions. However, the precision of direct sampling was still better. The artificial neural network (ANN) showed an improved accuracy and precision relative to the numerical gradient decent based approach in one dimension. However, when tested on real brain data, the ANN did not find any adequate solutions to the parameter estimations.

List of Abbreviations and Acronyms

MRI	Magnetic Resonance Imaging
SNR	Signal-to-Noise Ratio
SRR	Super-Resolution Reconstruction
SERO	Slice Excitation with Random Overlap
Slice Shifting	SRR method where one shifts the image a known subpixel amount with each image
Direct sampling	High-resolution conventional data sampling with isotropic voxels
Aspect ratio	The ratio between the through-plane resolution and the target resolution
DW-MRI	Diffusion-Weighted MRI
TR	Repetition time
TE	Echo time
T_1	Tissue dependent longitudinal relaxation time constant
T_2	Tissue dependent transversal relaxation time constant
b-value	Quantity describing the strength of the diffusion weighting
SE	Spin-Echo
PGSE	Pulsed-Gradient Spin-Echo
EPI	Echo-Planar Imaging
ANN	Artificial Neural Network

Contents

1	Introduction	1
2	Theory	2
2.1	Signal Generation	2
2.2	Relaxation Theory	2
2.2.1	Longitudinal Relaxation	2
2.2.2	Transversal Relaxation	2
2.3	Spin Echo with EPI Readout	4
2.4	Signal-to-Noise Ratio and Noise Floor Effects	4
2.5	Diffusion MRI	5
2.5.1	What is Diffusion?	5
2.5.2	Diffusion-Weighted Imaging	5
2.6	Super-Resolution Reconstruction	6
2.6.1	Acquiring the Low-Resolution Images	7
2.6.2	The Forward Sampling Model and the Optimisation Problem	7
3	Method and Experimental Setup	9
3.1	Design of the Experimental Sampling Matrix	9
3.2	Pulse Sequence Design	9
3.2.1	SERO Pulse Sequence	9
3.2.2	Diffusion Encoding	10
3.2.3	Optimisation of the Slice Selection Profile and Spoilers	10
3.2.4	Reconstruction of 2D-Images from Raw Data	11
3.3	Super-Resolution Reconstruction of Parameter Maps	11
3.3.1	Machine Learning Based Parameter Fitting	11
3.4	Numerical Simulations	13
3.4.1	Accuracy and Precision of Estimated Parameters	13
3.4.2	Reconstruction of a Numerical Phantom	14
3.5	Super-Resolution Reconstruction of In Vivo Data	14
3.5.1	Mapping Position Based Error Between Signal and High-Resolution Parameters	16
3.5.2	In Vivo Super-Resolution Reconstruction Using the Neural Network	16
4	Results	16
4.1	Numerical Simulations	16
4.1.1	Accuracy and Precision in Estimated Parameters	16
4.1.2	Reconstruction of Numerical Phantom	16
4.2	Demonstration of Super-Resolution Reconstruction by SERO In Vivo	19
4.2.1	Demonstration of the ANN In Vivo and Numerical Phantom	21
5	Discussion	22
5.1	The Benefits of SERO in Simulations for Accuracy and Precision	22
5.1.1	Reconstruction of Numerical Brain Phantom	24
5.2	Demonstration of SERO In Vivo	25
5.2.1	Demonstration of the ANN In Vivo	25
5.3	Conclusion and Future Aspects	26
A	Appendix	30

1 Introduction

Magnetic Resonance Imaging (MRI) stands as a cornerstone in modern medical diagnostics, offering unparalleled insights into the human body’s structural and functional characteristics [1]. Over the years, MRI has evolved, not merely as a tool for anatomical visualisation but as a modality for probing tissue microstructure and function at the molecular level [2, 3]. Through diffusion-weighted imaging it is possible to generate contrast based on the Brownian motion of water molecules, adding a new dimension to diagnostics through MRI based on microstructure rather than morphology [4].

The MRI image comprises both the acquired signal and random noise. In imaging human subjects, much of the random noise originates from the human body itself [5], leading to a bias in the signal that elevates its intensity [6]. This bias affects both the accuracy and precision of the experimental results. Hence a certain level of signal-to-noise ratio (SNR) is necessary to minimise the effect of this noise floor.

The signal bias is negligible at higher SNR, while at low SNR the accuracy of the experiment becomes greatly influenced by the bias. The SNR will increase by increasing the signal, and the impact of the signal bias will be reduced. Increasing the signal typically comes at an expense of spatial resolution. Another way to increase the SNR is to repeat the experiment and average the results. By doing this the precision of the experiment is improved [7] but the acquisition time increases with the number of averages.

Therefore, each MRI experiment is a trade-off between the image spatial resolution, the SNR and the acquisition time [7, 8]. Diffusion-MRI typically suffers from both poor resolution and poor SNR due to the nature of the experiment [9, 8]. A promising solution to this has in recent years been super-resolution reconstruction (SRR) methods [7]. These methods makes it possible to acquire data using thick slice selection profiles (i.e. with voxels that are larger in the through-plane direction relative to its in-plane direction) to increase the signal, leading to a higher SNR, while still recovering a high resolution in the through-plane direction.

In this study, super-resolution is enabled by acquiring multiple low-resolution diffusion-weighted images (where each image contributes with new information) that are used to directly reconstruct high-resolution diffusion parameter maps [8]. The purpose of this is to recover a high resolution parameter map while maintaining a high SNR. Typically this results in an underdetermined solution in a least-squares sense, making it ill-posed [7].

A novel sampling scheme is proposed with Slice Excitation with Random Overlap (SERO) instead of serial slice excitations, in order to collect the low-resolution images with variable intra-slice T_1 -weighting (as detailed in Section 2.6). Briefly, the idea behind this sampling scheme is to add a distribution of repetition-times (TR) instead of a constant time, adding another experimental parameter to the model, which improves the posedness of the reconstruction.

We implemented a reconstruction algorithm utilising an iterative gradient descent method to estimate diffusion and relaxation parameters from SERO-based data. Next, we assessed the precision and accuracy of the SERO method through numerical simulations on a digital phantom, comparing our results with similar methods documented in the literature [10]. We introduced a machine learning approach for parameter estimation, aimed at increasing the computational speed of the fitting process.

A research MRI pulse sequence tailored for SERO was developed using the Matlab-based Pulseseq framework [11], enabling practical experimentation at a human MRI scanner to demonstrate the functionality and efficacy of SERO through in vivo brain imaging, showcasing its potential in real-world applications.

2 Theory

2.1 Signal Generation

Magnetic resonance imaging (MRI) is based on the nuclear magnetic resonance (NMR) experiment on hydrogen nuclei. Each individual proton has a nuclear magnetic moment μ , which is proportional to its gyromagnetic ratio γ . For hydrogen nuclei, the gyromagnetic ratio is $\gamma = 42.6 \text{ MHz/T}$ [12]. In the absence of an external magnetic field, \mathbf{B}_0 , the spin orientation over a large quantity of spins is isotropic. However, in the presence of an external magnetic field \mathbf{B}_0 , the distribution of spins will slightly shift, resulting in a net magnetisation vector \mathbf{M}_z parallel to \mathbf{B}_0 . Each magnetic moment μ precesses about \mathbf{B}_0 with a frequency

$$\omega_L = -\gamma B_0, \quad (1)$$

known as the *Larmor frequency*. With a radio-frequency (RF) pulse \mathbf{B}_1 with frequency equal to the Larmor frequency ω_L , the net magnetisation vector \mathbf{M} will start to precess about \mathbf{B}_1 , and is in turn rotated down towards the xy-plane (given that \mathbf{B}_0 is parallel to the z -axis), where the \mathbf{B}_1 field is turned off (see Figure 1). In the xy-plane, a receiver coil is placed, where induction will occur due to the presence of a varying magnetic field (due to \mathbf{M}_{xy} precessing in the xy-plane). This will induce a current in the receiver coil, which in turn is the generated MR-signal.

2.2 Relaxation Theory

Once the RF-pulse is turned off, and $\mathbf{M} = \mathbf{M}_{xy}$, the system of in-phase spins contributing to \mathbf{M}_{xy} will start to gradually return to equilibrium and align itself along \mathbf{B}_0 . This gradual process is known as relaxation, and is divided into two main categories; longitudinal relaxation and transverse relaxation as seen in Figure 2.

2.2.1 Longitudinal Relaxation

In order for the system to return to equilibrium the abundant energy received by rotating \mathbf{M} down towards the xy-plane must be released to its surroundings. With longitudinal relaxation (also known as T_1 -relaxation), the excited nuclei return to its ground state (lower energy state) and re-align itself along \mathbf{B}_0 , resulting in a recovery of \mathbf{M}_z . This process is described by

$$\mathbf{M}_z(TR) = \mathbf{M}_0(1 - e^{-TR/T_1}), \quad (2)$$

where TR is the repetition time and T_1 the tissue/material dependent time constant known as the longitudinal relaxation time [13]. The longitudinal relaxation rate $R_1 = 1/T_1$, is enhanced by local magnetic field fluctuations with the surrounding environment of the spin, provided that the fluctuations occur at resonance frequency ω_L . This is less dominant for higher frequencies, implying that T_1 is prolonged for stronger B_0 due to its relationship to ω_L [14, 15].

2.2.2 Transversal Relaxation

The transverse relaxation, known as T_2 relaxation occurs due to the presence of local inhomogeneities $\Delta\mathbf{B}_0$ in the magnetic field \mathbf{B}_0 , which causes the spin system precessing in the xy-plane to experience different magnetic field strengths. This causes them to de-phase due to the difference in resonance frequency in regards to each other, resulting in the decrease of \mathbf{M}_{xy} (see Figure 1). Further, individual spins work like tiny magnets which in turn influence its neighbours, further causing

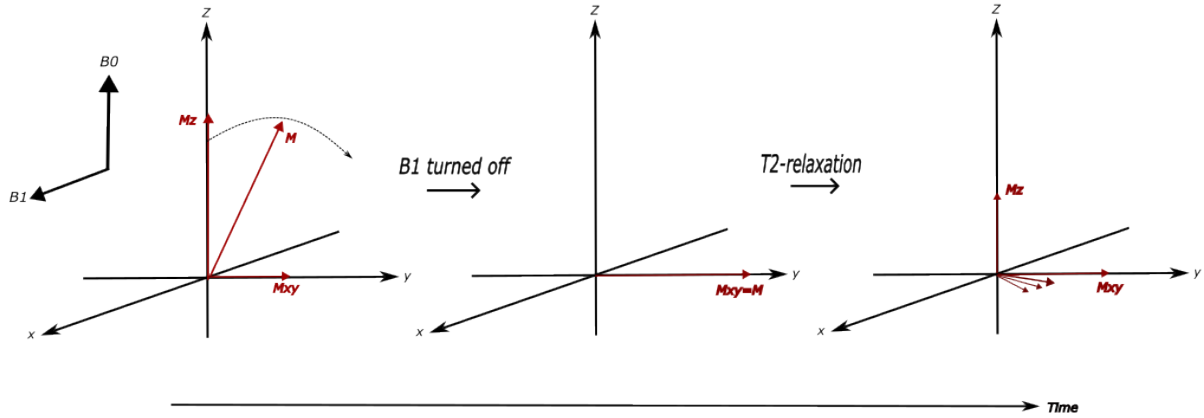


Figure 1 The de-phasing of the magnetic spins in the xy -plane, resulting in a decay of M_{xy} due to T_2 relaxation.

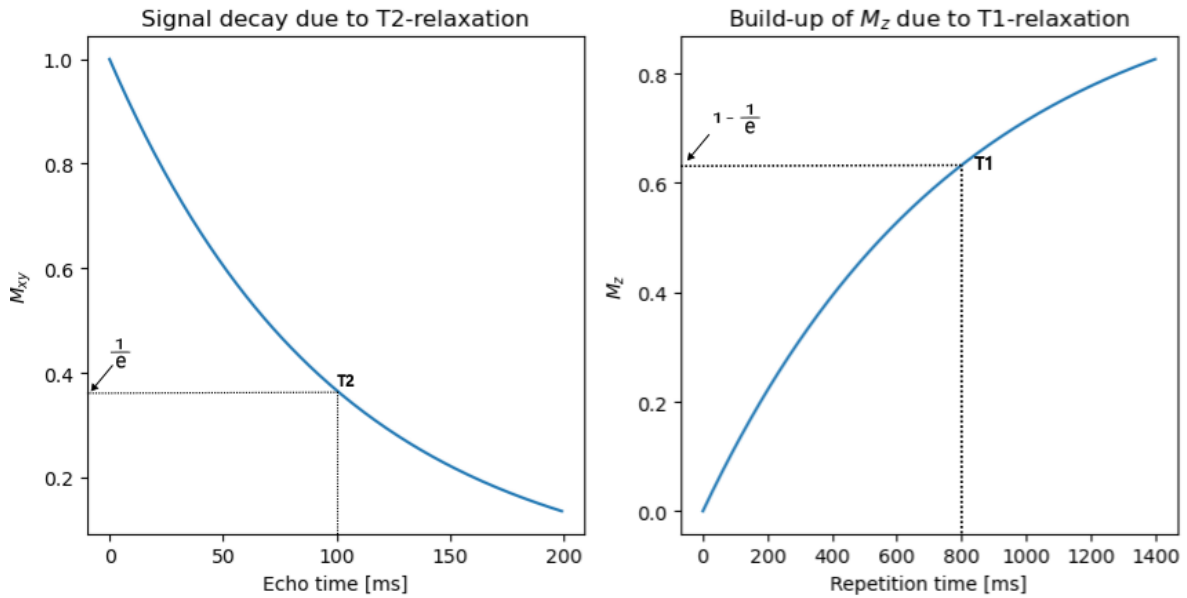


Figure 2 The two main relaxation types and their effect in the different components of the normalised magnetisation vector \mathbf{M} , for a tissue with a typical T_1 time of 800 ms and T_2 time of 100 ms.

local inhomogeneities in \mathbf{B}_0 . With a loss of phase coherence, the vector sum of the transverse magnetisation vector \mathbf{M}_{xy} decreases by

$$\mathbf{M}_{xy}(TE) = \mathbf{M}_0 e^{-TE/T_2}, \quad (3)$$

where TE is the echo time where the readout signal reaches its maximum and T_2 the tissue-dependent corresponding time constant known as the transverse relaxation time [13].




Figure 3 A sketch of a Spin Echo with EPI readout for fast imaging. The first 90° pulse represents the excitation pulse, and the 180° pulse the the refocusing pulse. The blue trapezoids represents the shifting frequency encoding gradient G_f , with the center of the echo represented as the red line.

2.3 Spin Echo with EPI Readout

To acquire any MR-signal, a pulse sequence is essential. It serves as a blueprint for the MRI-system, guiding the signal acquisition process. The spin echo (SE) pulse sequence begins with a 90° RF-pulse and slice selection gradient G_s . Then, a 180° RF-pulse, called a refocusing pulse, is applied at $TE/2$ along with the same slice selection gradient, reversing spin order and mitigating T_2^* effects. A frequency encoding gradient G_f is applied during signal readout, peaking at TE when spins are fully in phase before de-phasing again due to T_2 -relaxation. After a repetition time TR, the process is repeated with a new 90° RF-pulse.

Echo planar imaging (EPI) rapidly acquires MR-signals by alternating gradient polarities to generate multiple echoes, forming a zig-zag pattern in k-space. It introduces brief phase encoding gradient G_p activations after each echo to adjust spatial phase encoding, facilitating rapid signal acquisition. Despite its speed, EPI suffers from drawbacks like lower SNR, reduced resolution, susceptibility artifacts, and Nyquist ghosts due to rapid gradient polarity changes giving rise to eddy currents.

A regularly used pulse sequence for diffusion MRI is the single-shot EPI [16]. By integrating EPI into the SE framework as represented in Figure 3, rapid signal readout is achieved. This hybrid approach enables the assessment of diffusion properties across different anatomical regions. However, challenges such as susceptibility artifacts and reduced spatial resolution associated with EPI remains [17].

2.4 Signal-to-Noise Ratio and Noise Floor Effects

The inherent sensitivity of the MRI process makes it susceptible to various sources of noise, including thermal noise and artifacts. Hence, in MRI the signal-to-noise ratio (SNR) is a critical factor that influences the quality and reliability of the acquired images. The SNR represents the relationship between the strength of the signal from the desired tissue and the level of background noise. A higher SNR is desirable as it ensures visually clearer and more accurate images.

However, achieving a satisfactory SNR can be challenging in MRI due to several factors. One significant challenge arises from the trade-off between resolution and SNR. As the desire for higher spatial resolution increases, there is often a need to reduce voxel size or acquire thinner slices. Unfortunately, increasing the in-plane resolution or decreasing the slice thickness ΔZ leads to a lower amount of signal from the tissue of interest, negatively impacting SNR, as $SNR \propto \Delta Z$.

The summed MR-signal S is proportional to the number of spin excitations n per slice, whereas the noise will be proportional to the square root of n . Thus, $SNR = \frac{S}{N} \propto \frac{n}{\sqrt{n}} = \sqrt{n}$. Hence, by increasing the number of excitations of a slice, one can solve the SNR problem arising from the high-resolution thin slices. But this is of course at the cost of acquisition time, which is directly proportional to the number of excitations, n .

Moreover, the issue of SNR is compounded by the fact that diffusion-weighted MRI inherently involves weak signals. The MR-signals emanating from the human body are inherently faint, requiring sensitive detection equipment. Additionally, the presence of factors such as susceptibility effects, motion artifacts, and magnetic field inhomogeneities can further degrade the SNR.

In 1995, Gudbjartsson et al. provided crucial insights into the effects and characteristics of noise floor and signal biases [6]. These aspects had significant implications for the reliability and accuracy of MRI measurements. The noise floor represents the minimum detectable signal level in the MRI data. The noise arises from various sources such as thermal noise and electronic noise [5]. In regions with low SNR, the noise floor can obscure true signal changes, leading to uncertainties in the measurements. This in turn results in a signal bias when the complex MRI-signal is converted to a magnitude, which cannot be zero. Thus when noise is introduced, the magnitude operator is inflated in a positive direction. The signal bias refers to systematic errors or inaccuracies in the measured signal [6]. Consequently, accurate characterisation of tissue properties becomes challenging, impacting the validity of analyses and potentially compromising diagnostic accuracy. These biases distort the acquired signals, resulting in biased estimates of for example tissue parameters such as the diffusivity. As a consequence, inaccurate interpretations of tissue microstructure may occur, hindering the reliability of diffusion-MRI as a diagnostic tool.

2.5 Diffusion MRI

Diffusion-MRI is a specialised imaging technique that can provide information about microstructural organisation of tissue on a cellular level [2]. In other words it's a good way to probe tissue and infer quantitative information. In the prostate, for example, diffusion-weighted (DW) MRI can be used to accurately localise cancer growths [18], and has a potential ability in accessing the aggressiveness of the tumour [18].

In DW-MRI, the diffusion of water molecules is measured. Diffusion, referred to as Brownian motion, explains the random motion of particles in a medium, or random fluctuations in the particles positions. The motion of particles highly relies on the microenvironment of water, making DW-MRI reflect the abnormalities in different tissues [19].

2.5.1 What is Diffusion?

Diffusion corresponds to the random motions of particles in a medium. How quickly particles diffuse through a medium can be described by a diffusion coefficient D , quantifying the rate of self-diffusion. The diffusivity D can be described as

$$D = \frac{\langle r^2 \rangle}{6t}, \quad (4)$$

where r corresponds to the diffusion displacement in space and t is the time. The average displacement L as a function of time can thus be described as

$$L(t) = \sqrt{\langle r^2 \rangle} = \sqrt{6Dt}, \quad (5)$$

known as the Einstein equation for diffusion [20, 21].

2.5.2 Diffusion-Weighted Imaging

Typically for DW-MRI a spin echo with EPI readout is used [16], combined with two equally strong and long diffusion sensitisation gradients on each side of the refocusing pulse [17].

For quantification purposes, it is necessary to determine the degree of the diffusion-weighting. This is solved by summarising all gradient effects from both the diffusion- and imaging gradients into a term generally known as the *b-factor* [22, 23]

$$b = \gamma^2 \int_0^{TE} \left(\int_0^t G(t') dt' \right)^2 dt, \quad (6)$$

where γ is the gyromagnetic ratio and G the one-dimensional gradient waveform, which describes the spatial variation of the magnetic field. The inner integral describes the area of the gradient up to a time t , describing the spatial encoding imparted by the gradient during the time interval $[0, t]$. The outer integral captures the cumulative effect of the gradient waveform throughout the echo time.

The *b-factor* characterises the diffusion weighting of the sequence, in a similar way the TE characterises the degree of T_2 -weighting in a conventional spin-echo sequence.

With the *b-factor*, the signal attenuation from diffusion is reduced to a convenient and simple expression

$$S(b) = S_0 e^{-bD}. \quad (7)$$

This can be expanded by taking the Taylor expansion of the logarithm of $S(b)$ with respect to b , which gives a higher order term representing the variance V of the diffusion displacement, where

$$\ln \left(\frac{S(b)}{S_0} \right) = -bD + \frac{1}{2} b^2 V + \mathcal{O}(b^3), \quad (8)$$

and the higher order term $\mathcal{O}(b^3) \approx 0$. The variance term V becomes a measure of the diffusion displacement distribution, capturing the variability in diffusion behaviour, and is associated with the extent of diffusion heterogeneity in the medium [24, 25]. A higher variance thus implies a greater heterogeneity of the medium.

In DW-MRI the signal attenuation does not solely come from diffusion, but also influenced by T_2 - and T_1 -relaxation. It is thus important to account for signal attenuation not only due to diffusion, but also from relaxation effects (Equations 3 and 2). For a typical diffusion-weighted spin-echo sequence, the overall attenuation of the signal is given by

$$S = S_0 \left(1 - e^{-TR/T_1} \right) e^{-TE/T_2} e^{-bD + \frac{1}{2} b^2 V}. \quad (9)$$

2.6 Super-Resolution Reconstruction

In super-resolution reconstruction (SRR) the goal is to acquire a single high-resolution image from multiple low-resolution images. The necessity for SRR arises from the challenges MRI provides when it comes to resolution and SNR, as described in Section 2.4, due to the trade-off between resolution, SNR, and the acquisition time [7, 8].

Hence, new methods which enables the possibility to obtain a high-resolution image with acceptable SNR and scan time are currently being developed. This is where SRR is a favoured alternative as it allows reconstruction of a single high-resolution image from multiple low-resolution images with

a larger aspect ratio (i.e. the ratio between the through-plane resolution and the resolution we wish to achieve) [7] without any substantial increase in scan time. A specific acquisition scheme is combined with a specific SRR-method. The forward sampling model elucidates the process by which a high-resolution image of an object is derived from numerous low-resolution images of the identical object [8].

2.6.1 Acquiring the Low-Resolution Images

It is generally accepted that it is not possible to enhance the resolution in the in-plane dimension [8, 26]. Thus, low-resolution images are acquired with a high isotropic in-plane resolution, combined with a low-resolution through-plane dimension (i.e. thick slices) to improve the SNR of the acquired images. With thicker slices not only the SNR improves, the acquisition time shortens due to the fact that a larger part of the object is covered with each excitation. Finally, in order to achieve any enhanced resolution in the slice direction in the final high-resolution image, each low-resolution image must contain additional information about the object of interest. In order to achieve this, different strategies can be applied [8].

- **Slice Shifting:** One can acquire several low-resolution images as proposed by Greenspan et al. [10], where each image contributes with new information by shifting the image by a known subpixel amount with each excitation in the through-plane direction. In order to reach isotropic resolution of the voxel, a minimum of N low-resolution images are needed, with N being the ratio between the in-plane and through-plane resolutions. Slice shifting is an ill-posed problem and does not work well without a regularisation term in the optimisation problem [8, 27].
- **Slice Excitation with Random Overlap (SERO):** A novel approach to SRR was developed for this study by modifying the slice shifting approach to randomly shuffle the slice positions for each sample. By doing this over a large amount of repetitions, an entire object is covered, with varying TR per repetition, resulting in a distribution of TR. This gives SERO variable intra-slice T_1 -weighting. The extra experimental parameter will help the underdetermined nature of the system, making it less ill-posed.

In contrast to slice shifting or direct high-resolution sampling techniques, the SERO approach presents a unique advantage; inherent capability to provide T_1 -weighting. This is achieved through the distribution of TR inherent in the SERO design. Specifically, in SERO, the time intervals between overlapping slices are intentionally distributed randomly. This randomness contributes to the inherent T_1 -weighting effect. In contrast, in techniques like slice shifting and direct sampling, the time intervals between slices remain constant.

This distinction is visually represented in Figure 4, where the variable nature of TR in SERO is contrasted with the constant TR in slice shifting and direct sampling methods.

2.6.2 The Forward Sampling Model and the Optimisation Problem

The forward sampling model shows the relationship between the high-resolution image \mathbf{x} and low-resolution images \mathbf{y}_i . This can be represented as a transformation \mathbf{W}_i , mapping \mathbf{x} to \mathbf{y}_i ,

$$\mathbf{y}_i = \mathbf{W}_i \mathbf{x} + \mathbf{N}_i, \tag{10}$$

where \mathbf{N} represents the noise introduced during the imaging process. In practice, multiple low-resolution images $\mathbf{y} = (y_1, y_2, \dots, y_n)$ are acquired. Each is obtained with different transformations, capturing different information regarding the object at hand. The goal J is to find a high-resolution

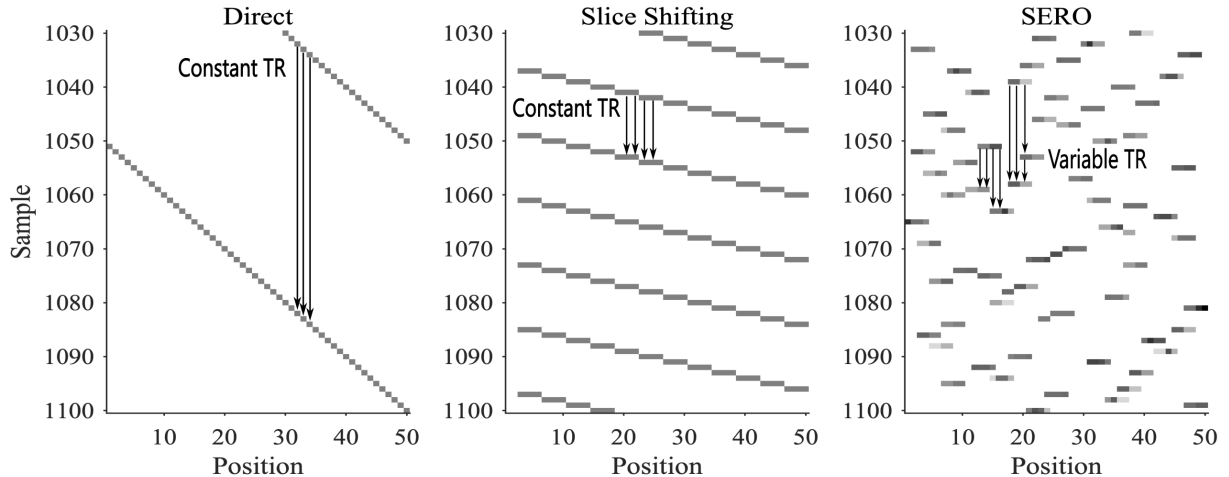


Figure 4 Sampling schemes for the SERO experiment with variable TR, and direct sampling and slice shifting with constant TR. On the x-axis the number of high-resolution positions is shown with thick slices distributed on them, and a subset [1030, 1100] of samples (where one sample is a repetition of the experiment) on the y-axis, in the much larger set of samples from [1, 2000]. The arrows in the figure represents the repetition time, i.e. the time it takes for an excited slice to be excited again. For direct sampling and slice shifting there is no varying greyscale intensity in the slices as appears for SERO, which shows that SERO has local differences in TR.

image \mathbf{x} that explains these low-resolution images. This is typically formulated as an inverse problem or optimisation problem, and can be solved using a least squares approach [7], i.e.

$$J(\mathbf{x}) = \sum_{i=1}^n \|\mathbf{y}_i - \mathbf{W}_i \mathbf{x}\|_2^2. \quad (11)$$

Typically, this results in an ill-posed problem in a sense that small deviations in the data can result in large deviations in the final solution of \mathbf{x} . To avoid overfitting to noise and enhance stability a regularisation term is necessary. This term penalises certain characteristics of the high-resolution image.

$$J(\mathbf{x}) = \sum_{i=1}^n \|\mathbf{y}_i - \mathbf{W}_i \mathbf{x}\|_2^2 + \lambda \|\mathbf{R}(\mathbf{x})\|_2^2. \quad (12)$$

Here, $\mathbf{R}(\mathbf{x})$ is a regularisation term, and λ a regularisation parameter controlling the trade-off between data fidelity and regularisation strength [7]. The point of the regularisation term is to contribute to shaping the desired properties of the high-resolution image. The optimisation problem seeks to find the high-resolution image $\hat{\mathbf{x}}$ that minimises the combined objective and the regularisation term

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \sum_{i=1}^n \|\mathbf{y}_i - \mathbf{W}_i \mathbf{x}\|_2^2 + \lambda \|\mathbf{R}(\mathbf{x})\|_2^2. \quad (13)$$

This can be solved using different optimisation algorithms such as the Levenberg-Marquardt algorithm for nonlinear equations or gradient decent [28]. These algorithms iteratively update the

estimate of \mathbf{x} based on the observed data and the regularisation term, and continues until a convergence criterion is met.

3 Method and Experimental Setup

In this study simulations were done to investigate the accuracy and precision of the novel sampling scheme for SRR. This was first done in one dimension by generating a bounded object containing 4 column vectors representing the parameters S_0 , the diffusivity D , T_1 and lastly the diffusional variance V , followed by reconstructions of a numerical brain phantom with the same parameters. The results were compared to slice shifting [10] and conventional high-resolution direct sampling.

In addition, in vivo measurements were performed using a research pulse sequence designed in the Pulseseq-interface [11], and the experiments were performed using a Siemens Magnetom Prisma MRI-system.

Lastly, an artificial neural network (ANN) was created in order to reduce the fitting time for the SRR described by Equation 13. The accuracy and precision of the ANN was compared to that of the Matlab least squares based *Lsqcurvefit*-function [29].

3.1 Design of the Experimental Sampling Matrix

The design of the experimental matrix was such that for each repetition of the experiment m , slices were excited randomly and shuffled over a given set of high-resolution positions n . In order to avoid too narrow overlap in certain areas, a set of historical repetitions were checked for possible overlap. If an overlap was detected within the given interval of repetitions, a new position was selected. For each repetition a b-value was used for diffusion-weighting factor from a selection of different b-values, with $b \in [100, 500, 900, 1400]$ s/mm².

The TR-matrix was designed by calculating the distance between a selected position and when it had been previously selected. The mapping W between the low-resolution plane to the high-resolution plane was calculated from the TR-matrix, where $\mathbf{W} = \mathbf{TR} > 0$.

Figure 5 shows a specific area of samples of the TR- and b-matrices together with its distributions. The TR-matrix visualise the position of the excited slices, where the intensity differences in the shuffled slices visualise local differences in TR.

3.2 Pulse Sequence Design

For the pulse sequence design, the Matlab-based Pulseseq framework was used [11]. The Pulseseq software framework is tailored for designing and simulating pulse sequences for MRI. It provides a comprehensive toolbox for constructing MRI sequences, including designing RF-pulses and gradient waveforms, assembling protocol designs and simulating their behaviour.

The parameters shown in Table 1 were used for the in vivo experiment.

3.2.1 SERO Pulse Sequence

A Pulsed Gradient Spin Echo (PGSE) example sequence provided with the Pulseseq library[30] with a single-shot spin echo with EPI readout was modified with SERO sampling, such that it would allow random slice positions. Position-wise repetition times and sample-wise b-values from SERO were used to create a sampling scheme suitable for random slice positions, including slice center positions and gradient/spoiler amplitudes for each sample.

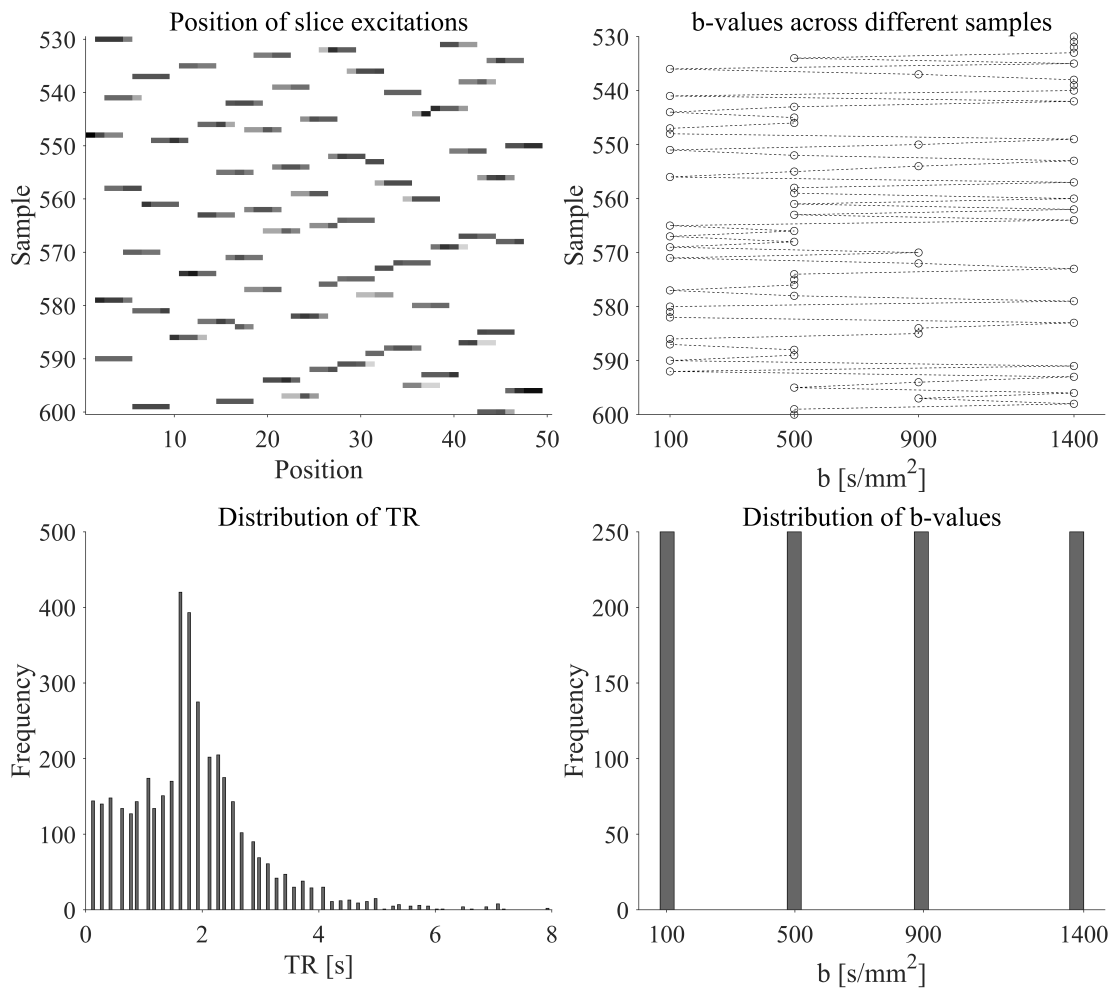


Figure 5 Top The experimental design of the SERO experiment, visualising the position of excited slices and the b-values across different samples (where each sample is a repetition of the experiment). Each voxel has an aspect ratio of 4, shuffled over $n = 50$ high-resolution positions for $m = 1000$ samples (where the top panels shows the interval $[530, 600]$), with a sample step time of 0.15 s, controlling the spacing of samples in time. The intensity differences in the slices visualises local TR, where a darker intensity represents a longer TR. The top right panel visualise the random shuffling of b-values across the different samples.

Bottom The distribution of TR- and b-values across $m = 1000$ samples.

3.2.2 Diffusion Encoding

The diffusion gradients were deployed in the (x, y, z) direction as trapezoids on each side of the 180° refocusing pulse. The amplitude of the diffusion gradients was scaled to achieve desired b-values to ensure precise diffusion encoding, with a maximal gradient amplitude of 49.3 mT/m.

3.2.3 Optimisation of the Slice Selection Profile and Spoilers

Modifications were added to the sequence by first improving the slice selection profile. This was by increasing the amplitude of the 90-degree RF-pulse. Its apodisation was changed from 0.3 to 0.5, changing its decay to zero intensity at the edges. Further, the time-bandwidth-product was increased from 4 to 10 in intensity, resulting in sharper edge transitions. To comply with hardware limitations the RF-pulse duration was increased from 3 ms to 6 ms (see Figure 6).

Table 1 Parameter values used for the pulse sequence.

Parameter	Value
Time between repetitions	0.15 s
Number of repetitions	1000
Number of positions	50
Aspect ratio	4
FOV	$180 \times 180 \times 75 \text{ mm}^3$
Acquired resolution	$1.5 \times 1.5 \times 6.0 \text{ mm}^3$
Target resolution	$1.5 \times 1.5 \times 1.5 \text{ mm}^3$
Slice thickness	6 mm
b-values	[100, 500, 900, 1400] s/mm ²
TE	80 ms
TR interval	0.02 – 21.9 s
Echo spacing	0.66 ms
Total acquisition time	150 s

These modifications to the RF-pulse resulted in a more box-shaped slice selection profile.

For regular sequences spoilers would not be necessary due to the long TR combined with plenty of T_2^* relaxation. However with SERO, the TR is typically short. Hence, the final modifications to be done was to add spoilers after each shot. This was done to avoid echoes from previous shots caused by the short TR. The resulting pulse sequence is shown in Figure 7.

3.2.4 Reconstruction of 2D-Images from Raw Data

The raw data was reconstructed by using a custom reconstruction program from Pulseq with included ghost correction and adaptive combine coil combination [31], with the raw data file and pulse sequence file as inputs. The k-space data was then based on the trajectory information from the pulse sequence file. The inverse Fast-Fourier-Transform was then used to convert the images from k-space to spatial coordinates, and saved into the final format representing the final volumetric image.

3.3 Super-Resolution Reconstruction of Parameter Maps

The SRR was performed using two different approaches. First through a least-squares approach by numerical gradient decent to solve Equation 13 using the Matlab-based Lsqcurvefit-function [29], a function used for nonlinear least squares curve-fitting. This was done by defining a forward signal model based on Equations 8 and 9,

$$\mathbf{S} \approx \mathbf{W} \odot \left(1 - e^{-\mathbf{TR} \oslash \mathbf{T}_1^T} \right) \odot e^{-\mathbf{bD}^T + \frac{1}{2} \mathbf{b}^2 \mathbf{V}^T} \mathbf{S}_0, \quad (14)$$

and using the designed 1000×50 \mathbf{TR} - and \mathbf{W} -matrices, and 1000×1 \mathbf{b} -matrix to solve for the experimental 50×1 parameter matrices $\mathbf{S}_0, \mathbf{D}, \mathbf{T}_1$ and \mathbf{V} , with $m = 1000$ number of samples, $n = 50$ numbers of high-resolution positions, and \odot, \oslash the *Hadamard product* and *Hadamard division* respectively.

3.3.1 Machine Learning Based Parameter Fitting

Due to the high computational time cost for numerical gradient decent approach to solve the SRR-problem, a machine learning based artificial neural network (ANN) was designed to solve this. It was developed using the Keras library [32].

For the training of the ANN, 10^6 normalised signal vectors each containing 1000 signals values were simulated using the forward model (Equation 14) at an SNR of 6. These were then split up into a training set and a test set, with the test set being 10% of the entire set of signal vectors, and the training set 90%.

To avoid excessive training time and overfitting, 20% of the training set was instead used for validation. The network with an architecture as described by Figure 8 was trained using the simulated signal vectors with a learning rate $\alpha = 10^{-5}$, and a batch size of 128. Each layer in the architecture except for the output layer used a ReLU activation function.

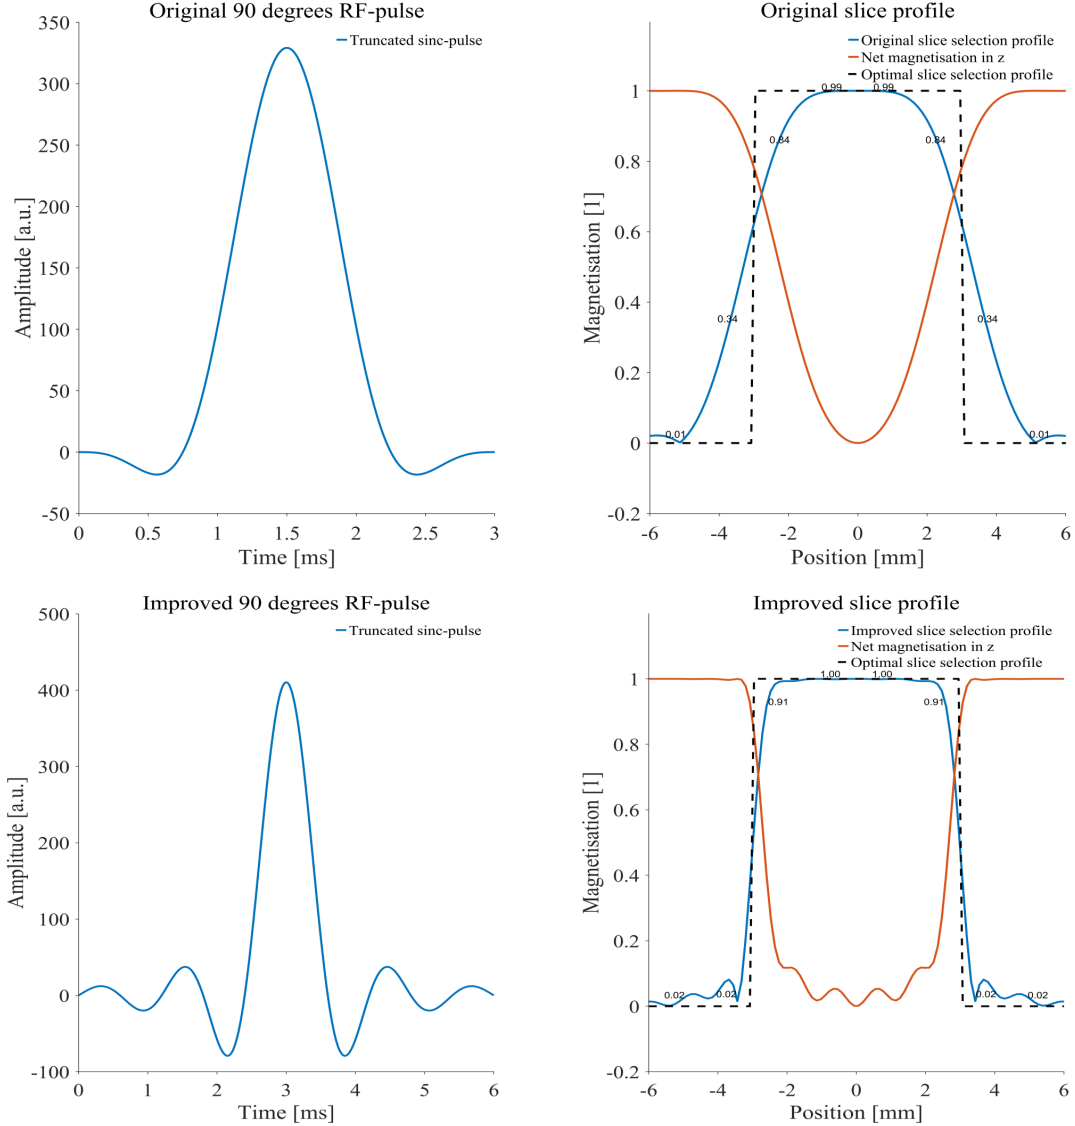


Figure 6 Improvement of slice selection profile by changing the shape of the RF-pulse, leading to a higher amplitude of the sinc-pulse. This resulted in a slice selection profile with a more box-like shape. In the original slice selection profile more than 30% of the available signal was outside the intended slice, whereas only about 2% in the improved slice profile was outside the intended slice.




Figure 7 The final design of the pulse sequence created using the Pulseseq interface, containing diffusion-encoding gradients, an EPI readout followed by crusher gradients.




Figure 8 The fully connected architecture of the ANN using an input layer of size 1000, followed by five hidden layers of sizes 2048, 1024, 1024, 512, 512 and an output layer of size 50×4 . The input to the ANN is a 1000×1 column of signal values from the through-plane direction of the low-resolution stack of images. The output from the ANN are the high-resolution 50×1 parameter maps S_0, D, T_1 and V .

3.4 Numerical Simulations

3.4.1 Accuracy and Precision of Estimated Parameters

A first test in order to see how SERO would perform compared to direct sampling and slice shifting was an analysis of accuracy and precision. Four random and bounded one-dimensional (1D) heterogeneous structures were generated containing 4 parameters and were contained in a matrix \mathbf{x} , where

$$\mathbf{x} = \begin{pmatrix} S_{0,1} & D_1 & T_{1,1} & V_1 \\ S_{0,2} & D_2 & T_{1,2} & V_2 \\ \vdots & \vdots & \vdots & \vdots \\ S_{0,n-1} & D_{n-1} & T_{1,n-1} & V_{n-1} \\ S_{0,n} & D_n & T_{1,n} & V_n \end{pmatrix}. \quad (15)$$

This was done in order to represent the ground truth value for the undecayed baseline signal $S_0 \in [0, 5]$, the diffusivity $D \in [0, 3.1]$, the longitudinal relaxation $T_1 \in [0, 1.5]$ and the variance $V \in [0, 0.8]$ distributed over $n = 50$ high-resolution positions.

The \mathbf{b} - and \mathbf{TR} -matrix was constructed and used together with the four 1D structures in order to

simulate a diffusion- and T_1 -weighted signal based on the forward model described in Equation 14.

This weighted signal was then used to test the SRR of the 1D object by using the forward model in Equation 14 to solve for the best possible fit $\hat{\mathbf{x}}$ in a least squares sense as described by Equation 13. This was performed in practice through both the Lsqcurvefit, using a gradient decent minimisation based solution to the SRR, and the developed ANN.

This process was repeated a set number of $N = 500$ times, where a new structure x_i was generated each time, resulting in a new fit \hat{x}_i . The accuracy of the experiment was estimated as the mean value of the differences between the ground truth and the fit for each experiment, i was calculated, i.e.

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^N (x_i - \hat{x}_i). \quad (16)$$

The precision of the experiment was given by the standard deviation σ of the error distribution arising from the simulations.

These numerical simulations were performed using SERO, slice shifting and direct sampling over an interval of SNR values, $SNR \in [1, 12]$, by applying noise to the input signal. The noise was introduced by combining two independent Gaussian noise components, each scaled appropriately according to the desired SNR, and calculating the resultant noise magnitude as the square root of the sum of their squares, resulting in a Rician noise distribution. Further, the noise was generated in a manner such that if the true signal is 1, the high-resolution image has an SNR equal to the desired SNR. Additionally, for a low-resolution image, the SNR is higher by a factor of the aspect ratio.

The accuracy and precision between SERO, slice shifting and direct sampling were compared to each other. Further, accuracy and precision for SERO between the Lsqcurvefit-based approach and ANN approach was compared, together with the time consumption to perform the SRR.

3.4.2 Reconstruction of a Numerical Phantom

To further evaluate SERO's performance reconstructions of a numerical brain phantom was done [33]. The three-dimensional (3D) brain phantom contained the same number of parameters (S_0 , D , T_1 and V) like the 1D structures did. The reconstruction of the numerical phantom was done in a similar way as described in the previous section for the 1D structure, by reconstructing the numerical phantom column by column in the low-resolution through-plane direction.

This was yet again done using SERO, slice shifting and direct sampling in order to see which method reconstructed the phantom in the most accurate way. The reconstruction was done using a perfect signal with infinite SNR as a reference reconstruction, followed by a reconstruction at an SNR of 6. In order to make a reference on how well the phantom was reconstructed, the letters LU were inscribed into the phantom in two places as seen in Figure 9.

3.5 Super-Resolution Reconstruction of In Vivo Data

With a pulse sequence compatible with random slice positions and the experimental design optimised for in vivo acquisition, a Siemens Magnetom Prisma MRI-system was used to collect the MRI raw data using a head coil. The signal from the raw data was reconstructed into the stack of the required transversal low-resolution images. These low-resolution images were used to reconstruct the high-resolution parameter maps described by Equation 13. A pipeline of the work is shown in Figure 10.




Figure 9 Numerical phantom of a brain with a spatial resolution of $0.1 \times 0.1 \times 0.1 \text{ mm}^3$ and with the letters *LU* inscribed with different sizes [33]. The pixels covered by the LU-letters had its mean signal intensity elevated by 40%. This phantom was then used as the ground truth when performing the reconstructions, and the purpose of the LU-markers was to assess how well the resolution was recovered.

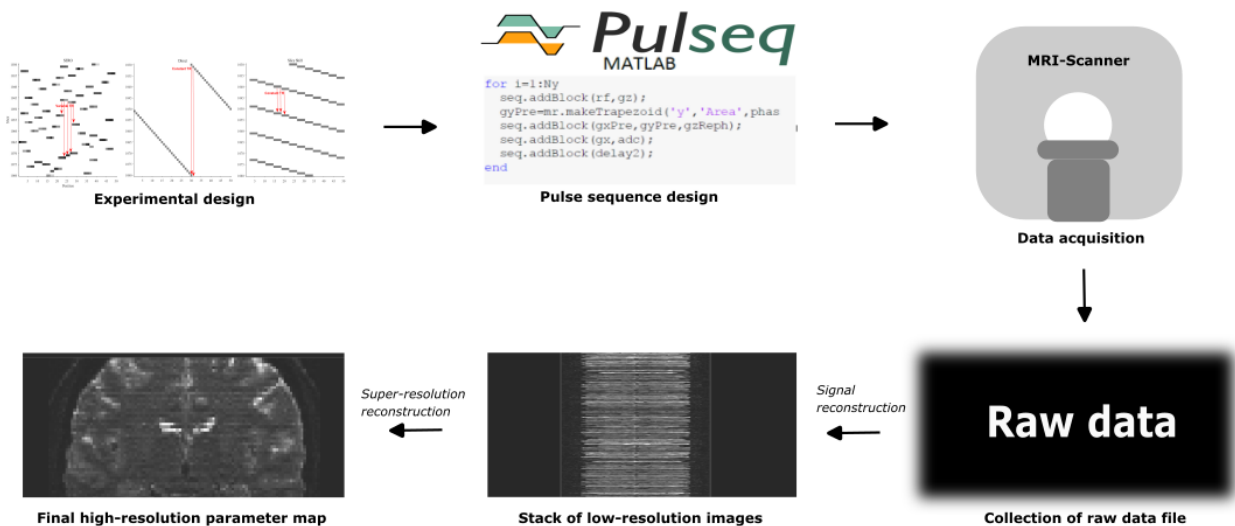


Figure 10 A pipeline of the work done to acquire the super-resolution reconstructed parameter maps from in vivo data. We started with an experimental design and sampling scheme which we implemented into a pulse sequence. The pulse sequence we could then bring to the MRI-scanner and collect our raw data. From the raw data we could reconstruct our stack of low-resolution images containing the signal vectors we would use to reconstruct the high-resolution parameter maps.

The data acquisition was done using the designed pulse sequence. One-thousand low-resolution transversal images were taken by using SERO with an aspect ratio of 1 (i.e. with already isotropic voxels) and 4 respectively. Another thousand transversal images were taken using high-resolution

direct sampling with an aspect ratio of 1. The raw data files were then reconstructed and converted into NIfTI images as described in Section 3.2.4.

Finally, from the low-resolution stack of transversal images, the high-resolution parameter maps were reconstructed using Lsqcurvefit.

3.5.1 Mapping Position Based Error Between Signal and High-Resolution Parameters

In order to understand where the SRR was not performing as expected, a position based error map of the brain was calculated, mapping the position-based signal difference between the signal given the forward model in Equation 14 and the acquired real data. The final error map was then constructed by taking the position-based weighted squared difference.

3.5.2 In Vivo Super-Resolution Reconstruction Using the Neural Network

The low-resolution transversal NIfTI images were super-resolution reconstructed using the ANN. The reconstruction was done by taking the signal values in the z -direction of the NIfTI-images column by column.

4 Results

4.1 Numerical Simulations

4.1.1 Accuracy and Precision in Estimated Parameters

The comparison of the accuracy and precision of SERO, slice shifting and direct sampling is shown in Figure 11.

The results show the bias induced at low SNR due to the noise floor effect. This has a strong effect on direct sampling compared to both SERO and slice shifting. The accuracy overall seem to be better with SERO compared to both slice shifting across all parameters except for S_0 . However, all sampling techniques seems to converge already at relatively low SNR, particularly for the diffusivity D . The accuracy for slice shifting in terms of S_0 appears to have a high error compared to both direct sampling and SERO.

In terms of precision direct sampling which has an aspect ratio of 1 has an advantage over SERO and slice shifting for S_0 , and a slighter advantage for the diffusivity. The precision seems also to converge as SNR increases, however at a much slower rate compared to the accuracy.

The comparison for accuracy and precision of SERO by the numerical gradient decent based Lsqcurvefit and the developed ANN is shown in Figure 12

From the simulations it appears that the ANN outperforms the Lsqcurvefit-based method in almost every aspect. It is only in terms of diffusivity that the accuracy is better for Lsqcurvefit. However, for the ANN the accuracy appears better for the remaining parameters, and the precision also improved for all parameters S_0, D, T_1 and V .

The difference in computational time for one fit was roughly a factor of 20 in favour of the ANN.

4.1.2 Reconstruction of Numerical Phantom

The phantom visible in Figure 9 was reconstructed using Lsqcurvefit, and results are shown in Figure 13 and Figure 14 respectively.

Figure 13 shows that with a signal in the absence of noise SERO reconstructs all 4 different parameter maps well. Both large and smaller LU-markers are clearly visible. Slice Shifting also

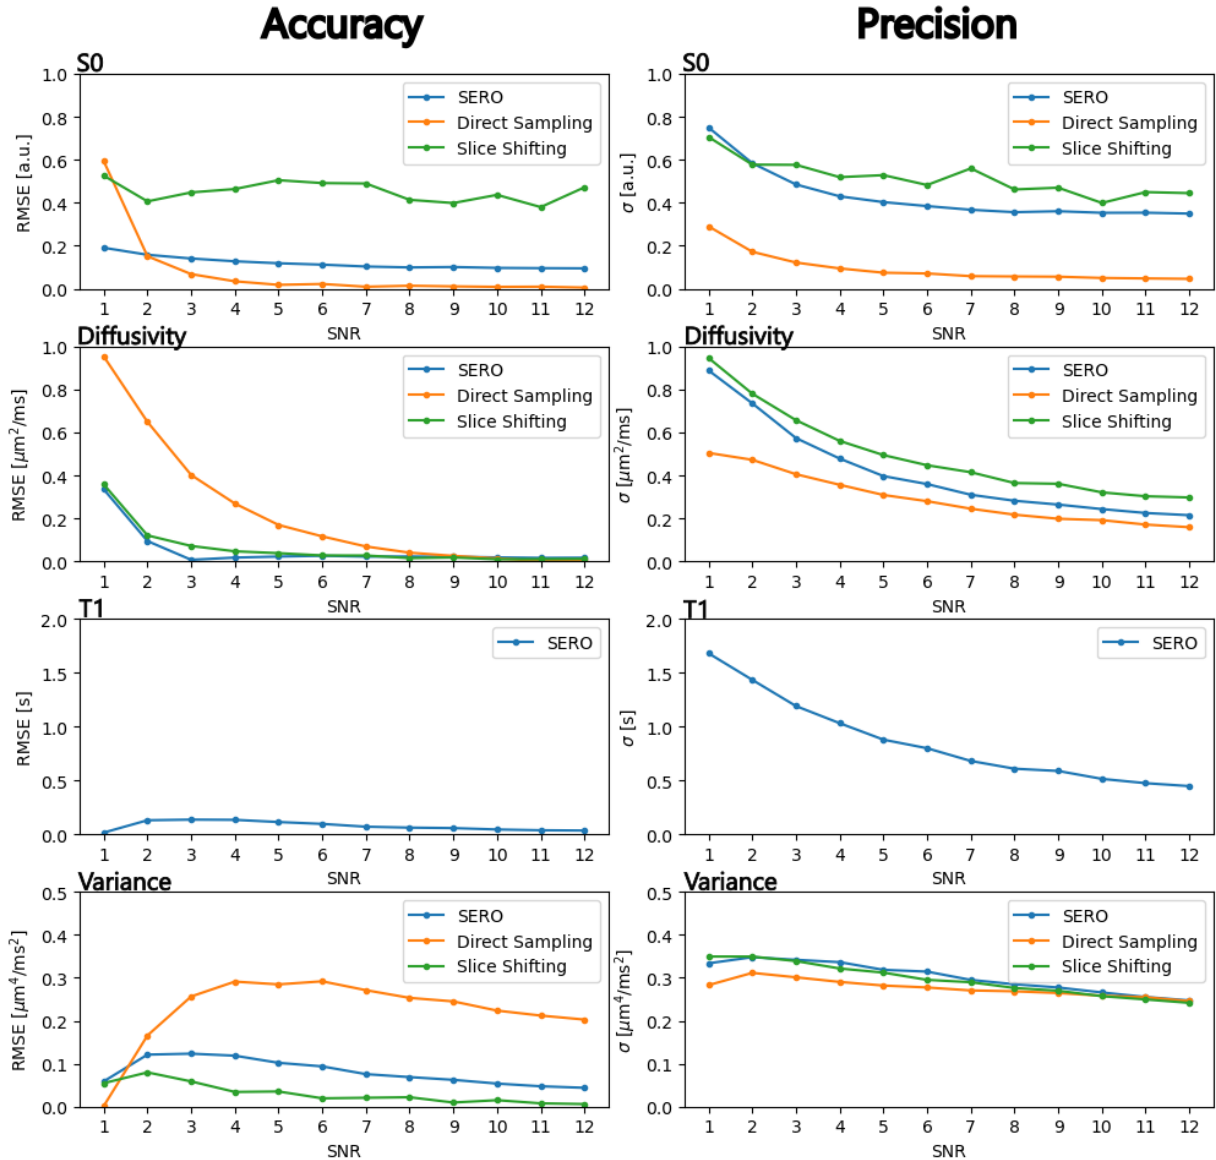


Figure 11 Numerical simulations for accuracy and precision of a 1D-structure over a range of SNR-values using the numerical gradient decent based Lsqcurvefit. The left panels shows the root mean squared errors (RMSE) from the simulations over a range of SNR-values, and the left panels the standard deviation σ of the error distributions from the simulations. Each panel visualise the differences between SERO, direct sampling and slice shifting in terms of accuracy and precision. Furthermore, the third row of panels visualising the accuracy and precision of the T_1 reconstructions only allows for SERO due to its inherent T_1 -weighting.

reconstructs the parameter maps with both LU-markers visible (except for T_1 due to it not having any inherent T_1 -weighting), although ring artifacts do appear as is more clearly shown in Figure 15. However, for direct sampling sampled at low resolution, the parameters are poorly reconstructed. Direct sampling fails to reconstruct the smaller LU-marker in all parameters and does not recover the resolution.

Figure 14 illustrates notable differences in reconstruction results. It shows that only SERO at a high resolution achieve satisfactory reconstruction of both the numerical phantom and the applied LU-markers for some of the parameter maps. Notably, even the small LU-marker remains visible in the S_0 parameter map. However, SERO encounters challenges in the other parameter maps, where

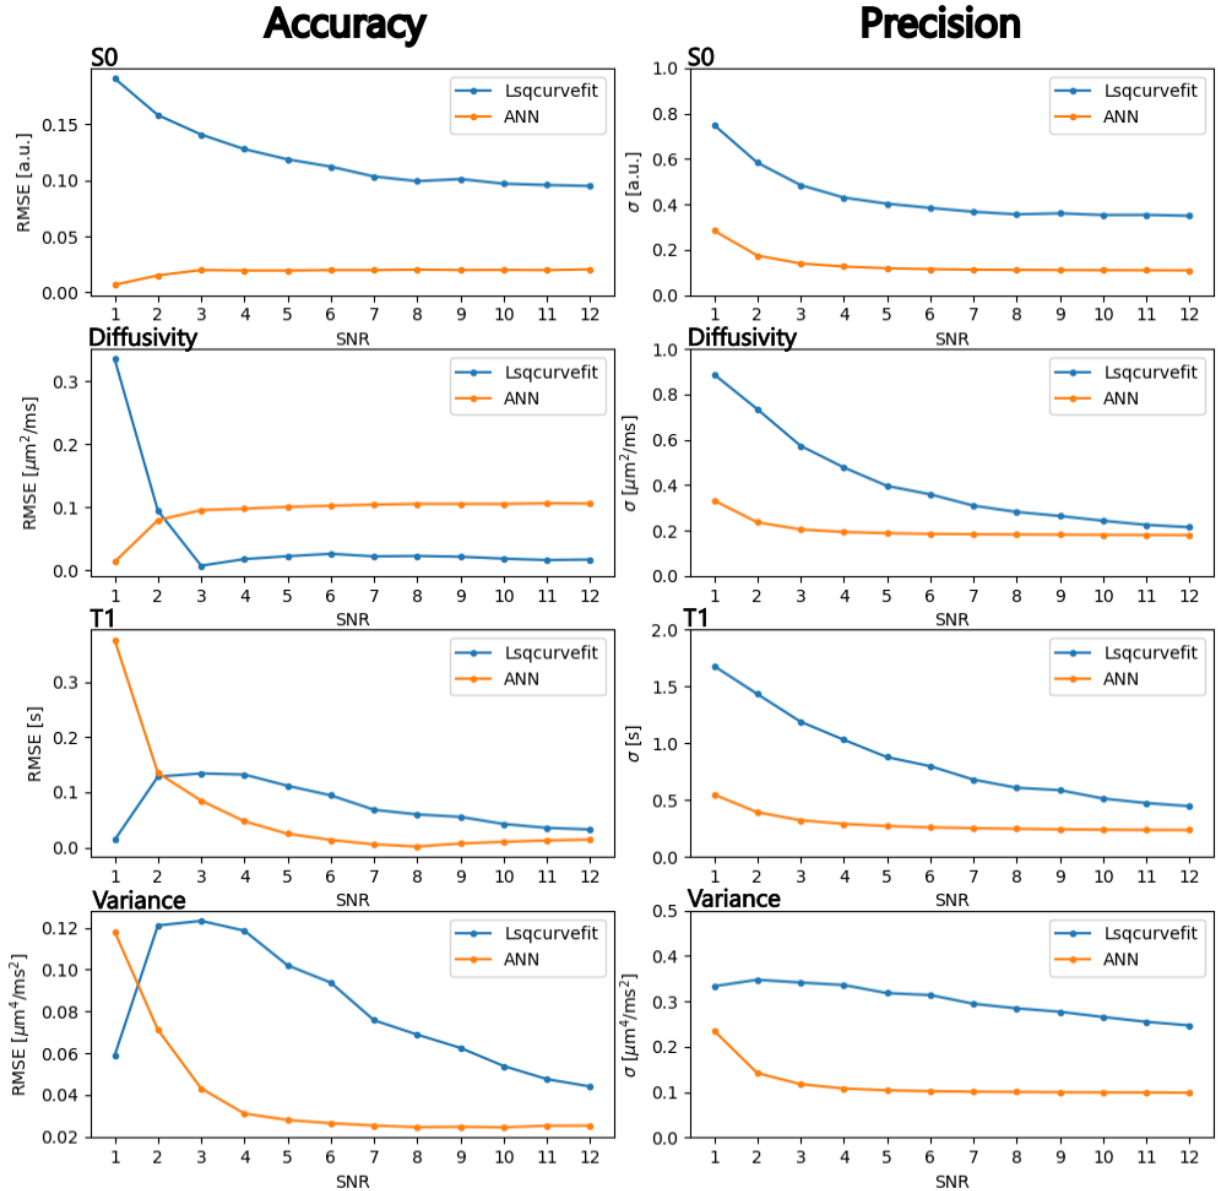


Figure 12 Comparison of accuracy and precision of a 1D-structure between Lsqcurvefit and the developed ANN. We see an improved accuracy and precision for the ANN compared to Lsqcurvefit. Particularly the precision is improved for the ANN, with the biggest impact in the lower SNR-regions. Thus, the ANN provides a higher accuracy and precision compared to the numerical gradient decent based approach, while also increasing the computational time by a factor of 20.

the larger LU-marker remains discernible only in the diffusivity map but disappears in the T_1 and variance maps due to poor contrast. Additionally, the smaller LU-marker becomes indiscernible in these three maps.

Slice Shifting barely manages to adequately reconstruct the larger LU-marker in the S_0 map, with the smaller LU-marker being lost. Moreover, in the diffusivity map, even the larger LU-marker is poorly reconstructed, and the super-resolution reconstructions appear unsuccessful in all three parameter maps.

High-resolution direct sampling seems to recover resolution in the S_0 map well, where both large and small LU-markers are visible. However, the LU-markers in the other parameter maps are

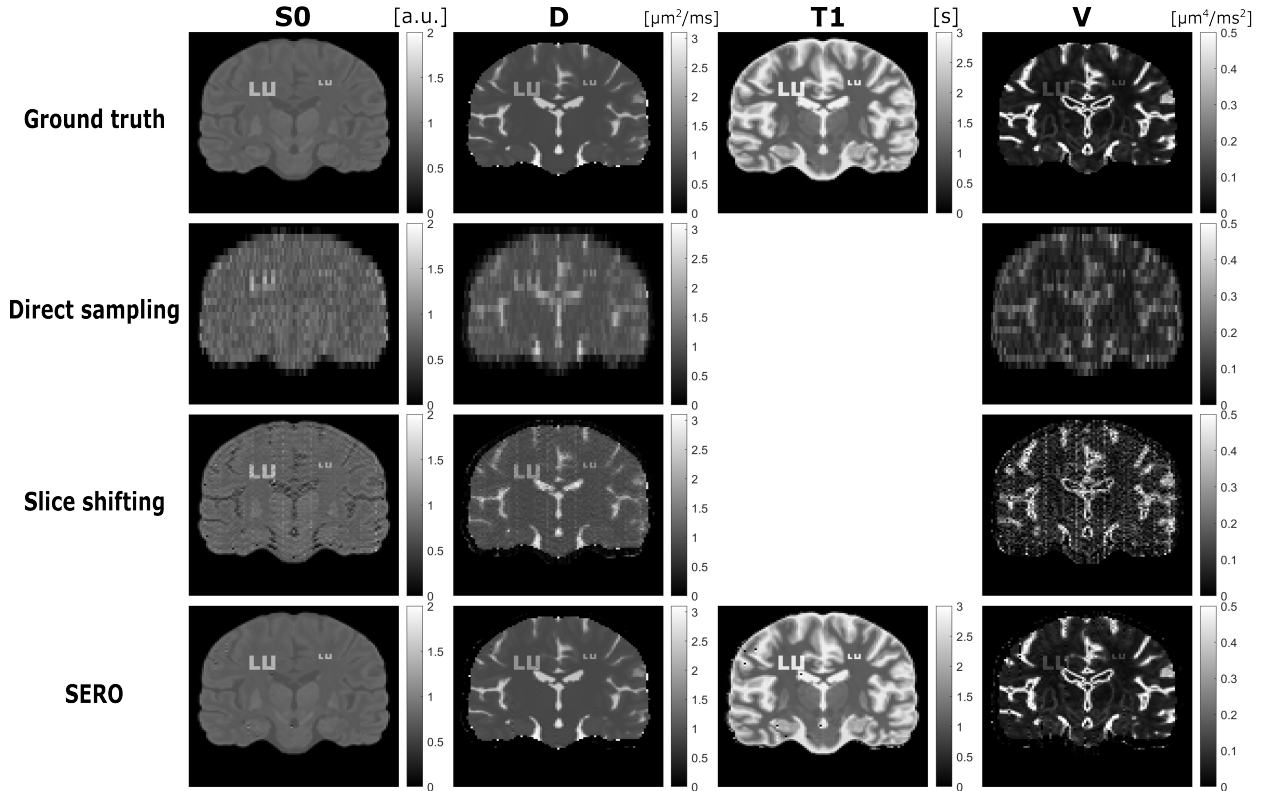


Figure 13 Super-resolution reconstruction of the numerical phantom by Lsqcurvefit given a signal at infinite SNR, acquired with an aspect ratio of 4 using direct sampling, slice shifting and SERO. SERO appears to more than adequately recover the high resolution, whereas no high resolution is recovered using direct sampling. The high resolution is recovered with slice shifting. However, ringing artifacts appear, and are particularly visible in the S_0 map.

obscured by noise. Conversely, the low-resolution direct sampling case handles noise better, but fails to recover high resolution in any of the parameter maps, rendering the smaller LU-markers invisible.

4.2 Demonstration of Super-Resolution Reconstruction by SERO In Vivo

In Figure 16 the SRR in vivo data of a brain is shown.

The direct sampling manages quite well to reconstruct the S_0 map and diffusivity map. However, it appears to work rather poorly when reconstructing the diffusional variance map. The contrast in these maps is quite poor, and the reconstruction appears to be very noisy.

When using SERO at high-resolution (i.e. with an aspect ratio of 1), we see signal artifacts or inaccuracies in the signal model in the form of the horizontal streakiness in the images. And similar to the case when direct sampling is applied, the contrast in the diffusional variance map is rather poor.

Using SERO at low-resolution with an aspect ratio of 4 seems to resolve some of the issues observed with the other approaches. The three parameter maps S_0 , D and T_1 are reconstructed fairly adequately with a descent contrast. However, the reconstruction of diffusional variance map V appears to fail in the peripheral zones, resulting in an uneven and streaky parameter map. Furthermore, the intensity range for the S_0 map appears much larger when using SERO with an aspect ratio of 4 than for the other two methods.

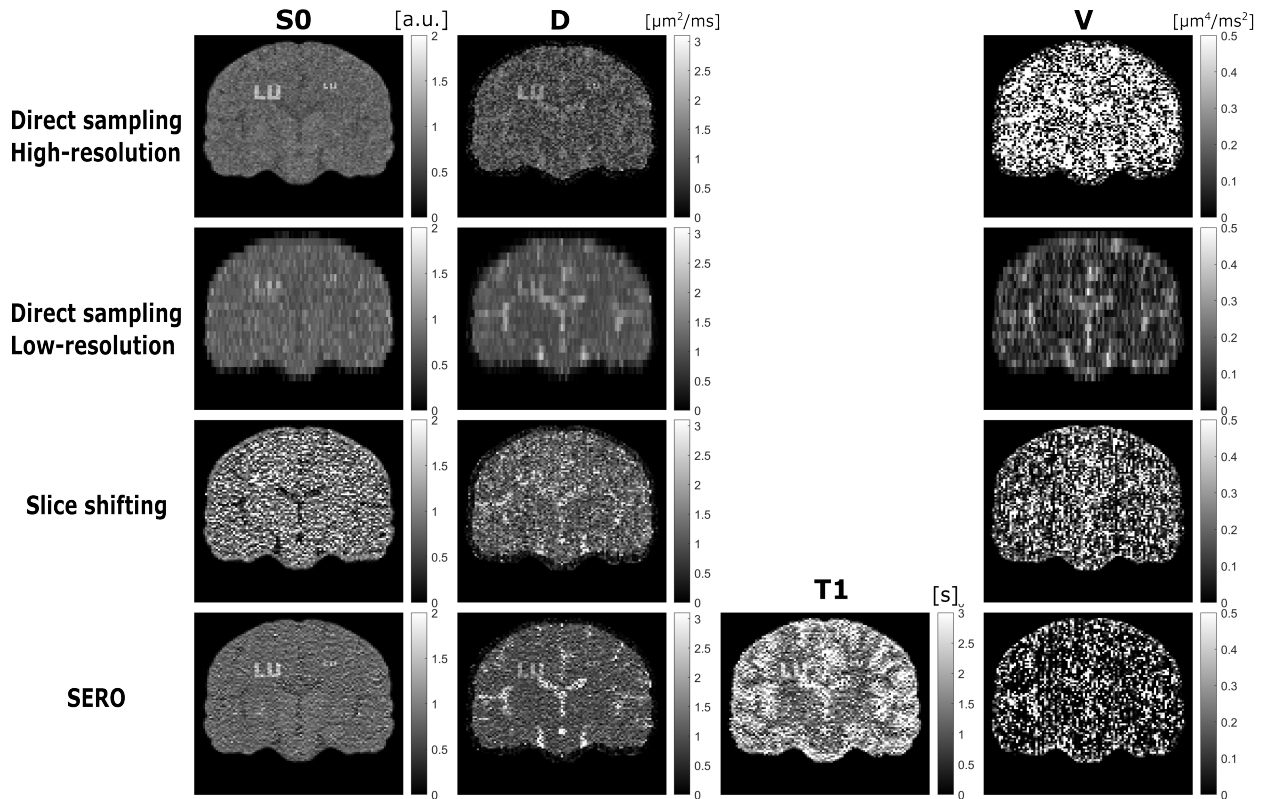


Figure 14 Super-resolution reconstruction of numerical phantom by Lsqcurvefit given a signal with a rectified noise at an SNR of 6. Direct sampling with high resolution was reconstructed with isotropic voxels. The remaining images were acquired using an aspect ratio of 4.

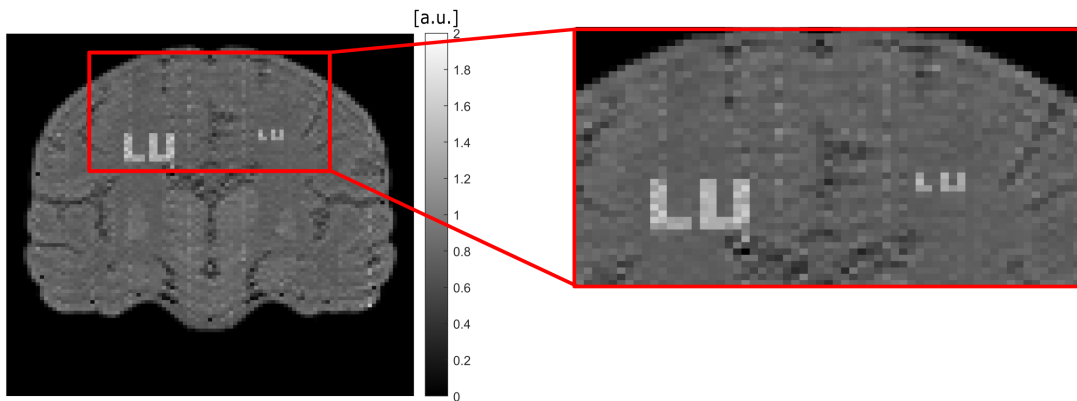


Figure 15 Periodic ringing artifacts arising when reconstructing the numerical brain phantom using Slice Shifting with an aspect ratio of 4 in the S_0 parameter map.

Figure 17 shows a transversal view of the spatial signal errors, displaying in what areas the reconstruction method was better and what areas were worse in the 5th, middle and 45th slice of the brain image.

It appears that the position based errors are larger in the peripheral zones compared to more central zones. Furthermore, areas of transition between different tissues appears to have larger deviations. An example of this is the transition between the parenchyma and the cerebrospinal fluid, or between white-matter and the ventricles. Additionally, there are relatively large errors in the ventricles. Also, ghosts from the EPI-readout appears visible in the middle slice.

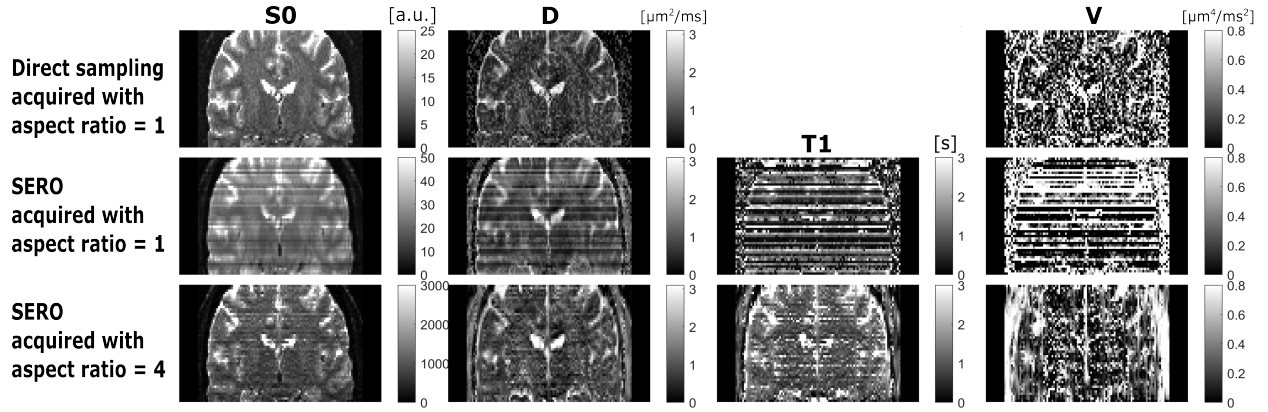


Figure 16 SRR of in vivo brain data with high-resolution direct sampling and SERO, acquired using readily isotropic voxels with an aspect ratio of 1, followed by anisotropic voxels with an aspect ratio of 4 (see appendix A). When using direct sampling, the parameter maps appear to have been reconstructed adequately (except for the diffusional variance map) without any clear regions where the reconstruction fails. However, for the reconstruction by SERO acquired with an aspect ratio of 4, the peripheral zones of the diffusional variance map appears to have failed. Furthermore, the reconstructions by SERO acquired with an aspect ratio of 1 has an inherent horizontal streakiness in every parameter map.




Figure 17 Signal and position based error map of the SRR. Areas with a large deviation appears brighter. The red arrows in the left most panel points toward the peripheral areas where larger errors were typically found. Furthermore, in the center panel the arrows points towards a Nyquist ghost and the cerebrospinal fluid in the ventricles. In the right panel we see large signal errors again in the ventricles and also in the eyes.

4.2.1 Demonstration of the ANN In Vivo and Numerical Phantom

In Figure 18 the SRR with Lsqcurvefit is compared with the SRR through the developed ANN based on the same input data.

It appears that the ANN does not seem to successfully reconstruct the image based on the low-resolution input data. In the transversal plane it manages to reconstruct the in vivo data somewhat adequately. However, it is clear that in the coronal- and sagittal planes, the reconstruction fails and nothing valuable is reconstructed. The ANN does however reconstruct the numerical brain phantom in all three planes.




Figure 18 Super-resolution reconstruction of the S_0 parameter maps of in vivo data through the numerical gradient decent based Lsqcurvefit, and the self developed ANN. The bottom panel shows the super-resolution reconstruction of the numerical brain phantom using the ANN to demonstrate that it works to some extent, even though there are clear distortions and artifacts in the images. The phantom reconstruction exhibits streaky behaviour in the coronal and sagittal planes, while in the transverse plane, multiple slices appear to overlap, distorting the image. In the central panels, the reconstruction of the in vivo data using the ANN fails.

5 Discussion

5.1 The Benefits of SERO in Simulations for Accuracy and Precision

From numerical simulations for accuracy and precision we found that when compared to conventional high-resolution direct sampling and the Slice Shifting method, SERO had the overall best accuracy, while direct sampling has a higher precision.

Figure 11 demonstrates that the accuracy of SERO is superior compared to slice shifting and direct sampling. For S_0 , slice shifting struggles to find a unique solution to Equation 13, even at higher SNR levels. This difficulty is further reflected in the precision metrics, where slice shifting shows minimal improvement compared to SERO and direct sampling.

The precision for direct sampling is better compared to both SERO and slice shifting in every case except for T_1 . This is expected since when using direct sampling, the voxels are already in an isotropic state, and does not involve SRR and does not suffer from inflated variance due to the ill-posedness of the inversion process.

From the simulations it is also visible that direct sampling struggles at lower SNR, where the noise floor effect has a higher impact, resulting in a larger bias. The RMSE for direct sampling is inflated by the signal bias for S_0 , D and V . However, the signal bias effect appear stronger for the diffusivity and diffusional variance compared to S_0 . The strong inflation of the RMSE of direct sampling compared to SERO and slice shifting is an expected result as the slices are acquired using high-resolution voxels, resulting in a decrease in SNR, resulting in a larger bias from the noise floor. SERO and slice shifting seems to be excelling in these areas, and are not as affected of the noise

floor. Similar results have been seen in previous studies where SRR is shown to improve accuracy [34] and spatial resolution [8, 7] in the through-plane direction in comparison to direct sampling.

We note that SERO reaches a high accuracy at relatively low SNR in all four parameter maps (with a relative bias of less than 10% at $\text{SNR} > 7$). Direct sampling on the other hand has a relative bias of less than 10% already at an $\text{SNR} > 3$ for the S_0 map. For the diffusivity map it does not reach this relative bias until an $\text{SNR} > 8$, and for the diffusional variance map, the relative bias for direct sampling does not become less than 20% within the scope of the simulations. This shows how much the signal bias inflates the signal intensity when direct sampling is used. The RMSE from SERO and direct sampling converges at an $\text{SNR} \approx 8$ for the diffusivity maps, and already at $\text{SNR} = 2$ for the S_0 map. Hence, it appears that SERO only has a true impact at the really low SNR regions of $\text{SNR} < 8$ for the diffusivity, and $\text{SNR} < 2$ for S_0 . Only for the diffusional variance map does SERO have a large impact at $\text{SNR} > 12$ relative to direct sampling.

Further, in terms of accuracy, based on Equation 14, S_0 should be the least complicated parameter to find a fit for, and the variance the most complex. This could explain why the accuracy for direct sampling for S_0 is so good even at $\text{SNR} > 2$ compared to SERO. These SNR values appear low, so it is worth noting that the noise was added in a manner that if the true signal was 1, the SNR would be true. However, in these situations the true SNR of the parameter maps would be somewhat elevated.

Throughout our experiments, we made qualitative observations that the estimation of T_1 was more challenging than S_0 and the diffusivity maps. The T_1 maps generally had poor precision and visual quality. However, from the numerical simulations the accuracy any precision increases as the SNR increases as we would expect, and the effect of the noise floor is minimised. It is also worth mentioning that the very small bias at $\text{SNR} = 1$ comes with a very large standard deviation, meaning that the distributions of errors after the 500 simulations at this SNR was wide. Hence, it only appears that the bias is small at an $\text{SNR} = 1$, and then increases and decreases again. This is an inaccurate behaviour and may be an effect of the fitting procedure at extremely low SNR (the RMSE reaches a maximum at $\text{SNR} \approx 3$, and then starts to decrease again as the SNR increases).

In summary, SERO demonstrates superior accuracy compared to direct sampling at lower SNR, aligning with the findings of Vis et al. [34], which showed that SRR improves accuracy over direct sampling. However, the precision of SERO was found to be somewhat lower. The precision of SERO could potentially be enhanced through the use of regularisation in SRR. This would however result in a decrease in accuracy. Therefore, a balance between regularisation for high precision and maintaining accuracy needs to be established.

Lastly, the developed artificial neural network (ANN) was tested against the numerical gradient decent based Lsqcurvefit-function, also in terms of accuracy and precision of SERO. The ANN was found to be superior both in terms of accuracy and precision, but especially in terms of computational time. The time for the SRR was decreased by a factor of 20 with a trained ANN. This meant that the same 500 simulations could be performed in a matter of minutes instead of hours.

The precision achieved by the ANN surpasses that of Lsqcurvefit, as anticipated. Regression-based artificial neural networks often come with regularisation due to their architecture and training process, where both weight decay (i.e. by adding a penalty term to the loss function that discourages large weight values, preventing overfitting), early stopping (where the training is stopped early based on the validation of the training) and batch normalisation (used to improve training stability and speed, where the activations of each layer is normalised), which all can act as a form of regularisation by reducing internal covariate shift and helping prevent the model from overfitting. These regularisation techniques are often built into the training process of neural networks for regression tasks to improve generalisation performance and prevent overfitting, making them

inherently regularised. This is also what we observe in terms of the precision being superior for the ANN.

The most significant advantage of using the ANN is the enhanced computational speed for performing SRR. Even this simple regression-based ANN already outperforms the nonlinear least squares solution in terms of both accuracy and precision, while also being much faster. However, it is important to note that the ANN is less robust to deviations from the training data.

5.1.1 Reconstruction of Numerical Brain Phantom

SRR by SERO, slice shifting and direct sampling was performed in a similar way as in the one-dimensional case. However, instead of generating a heterogeneous 1D-object, a numerical phantom of a brain was used which contained the same 4 parameter maps S_0, D, T_1 and V . It was found that SERO reconstructed all 4 parameter maps well with a signal at infinite SNR, whereas slice shifting recovered the resolution, but ringing artifacts emerged. Furthermore, when using direct sampling at a low resolution no resolution could be recovered in the through-plane direction. When the signal from the brain phantom was contaminated by noise at SNR= 6, SERO still managed to reconstruct all 4 parameter maps, however, only S_0 and D adequately. The reconstructions of the T_1 - and V maps were too contaminated by noise.

In Figure 13 we see the numerical phantom reconstructed using a signal at infinite SNR based on Equation 14. Here we see that direct sampling at low resolution encounters difficulty when trying to reconstruct the object for the parameter maps. It does succeed to reconstruct the phantom on a larger scale, but for smaller objects in the phantom (such as the LU-markers placed on two locations with different sizes), it does not even succeed to adequately reconstruct the larger accuracy marker in the S_0 case, which clearly is not any problem for either SERO or slice shifting.

In order to perform any fit, a distribution of minimum two parameter values are needed, and since direct sampling and slice shifting has a constant TR, no fit is possible for the T_1 maps for direct sampling and slice shifting. In contrast, SERO manages successfully to reconstruct all 4 parameters adequately. Further, it is interesting to note that ring artefacts appear when using slice shifting. As seen in Figure 15, these ringing artifacts are periodic and based on the aspect ratio used when acquiring the signal. These show the underdetermined nature and the ill-posedness of the SRR problem when using the slice shifting technique, which was also exhibited by Nicastro et al. [27]. These ringing artifacts do not appear when using SERO, which demonstrates that the posedness of the SRR problem is improved by its inherent T_1 -weighting. Furthermore, the reconstructions align with the conclusion of Plenge et al. [7], that the trade-off between spatial resolution, SNR and acquisition time can be improved by SRR.

A notable contrast emerges when reconstructing the phantom using a signal contaminated by noise. The reconstruction by direct sampling at a low resolution results in the parameter maps having poor precision, and does not allow the identification of the brain morphology. Direct sampling at high resolution successfully reconstructs both LU-markers in the S_0 map, as expected from the numerical simulations even at SNR= 6. However, the remaining maps are contaminated by noise, resulting in a poor reconstruction. This is also in terms with what we see from the 1D simulations, where direct sampling struggles in terms of accuracy at SNR= 6, where it is more affected by the noise floor. However, slice shifting also appears to fail at this SNR. Both the diffusivity- and variance maps appears almost homogeneous with noise, and it is clear that the reconstruction has failed. The S_0 map shows the general features and the larger accuracy marker (although contaminated by noise). Neither of the LU-markers appears in the remaining maps, and are lost in the noise of the images. SERO on the other hand succeeds in this situation. In the S_0 case even the small accuracy marker is still reconstructed adequately, and is clearly visible. The small marker is lost in the noise in the diffusivity map, but the large accuracy marker is still visible. However, for the T_1 map, the larger LU-marker is barely visible in the noise, whereas the smaller LU-marker is lost due to poor

contrast. For the diffusional variance map, both markers are lost due to poor contrast.

In summary, the parameter maps appear to comply relatively well for direct sampling and SERO with what was seen in the 1D-simulations. From Figure 11 we expect that direct sampling at high resolution can only reconstruct the S_0 map with high accuracy at SNR= 6, which is also what we observe. The remaining parameter maps are however heavily biased by the noise floor, which aligns with what was shown by Van Steenkiste et al. [8]. SERO on the other hand manages both situations (i.e. when using a signal at infinite SNR, and a signal at SNR= 6) adequately, as would have been expected from the 1D-simulations. However, slice shifting should perform better at the lower SNR regions based on the 1D-simulations, except for S_0 where the accuracy is low. It is clear from Figure 14 that this is not the case. It would also be expected to see more of the underdetermined nature of slice shifting in form of a horizontal streakiness in the images. This streaky behaviour is however not visible either.

5.2 Demonstration of SERO In Vivo

The SRR was done on in vivo data by SERO with an aspect ratio of 4, and compared with direct sampling and with with SERO using an aspect ratio of 1. SERO with an aspect ratio of 4, along with direct sampling managed to reconstruct the S_0 - and diffusivity maps adequately. From Figure 16 we see that both direct sampling and SERO with an aspect ratio of 4 finds a good fit to the solution for the S_0 map. This complies with what was seen in the 1D-simulations. In the case of the acquired brain data, the SNR was expected to be generally high. Based on the 1D-simulations, both SERO and direct sampling would perform well at high SNR, both in terms of accuracy and precision, and that is what corroborated by the obtained results, in for example the S_0 map. For all 4 parameter maps we do see the expected streaky behaviour in the images for SERO with an aspect ratio of 1, hinting at the underdetermined nature of the SRR-problem.

In the error map visible in Figure 17 we see what areas were more poorly reconstructed and which areas were better. We see that there has been some ghosting likely arising from the EPI-readout [35]. This has in turn affected the signal inside the brain, and thus also the final parameter fit, which takes in the signal values as inputs. Also it is clear that in the areas of transition between different tissues the deviations are larger. An example of this is the transition between white- and grey-matter, or between white-matter and the ventricles. This is likely due to the sharp change in parameter values in these regions, making it harder to find a suitable fit.

5.2.1 Demonstration of the ANN In Vivo

The ANN was used to reconstruct in vivo data of a brain, which was compared to the same data reconstructed by a least squares approach. From Figure 18 we see that the ANN has not worked on the in vivo brain data. Based on the low-resolution input data, it does not manage to find any good solution to the SRR. The transversal plane is reconstructed adequately in order to see the shape of the brain, which would be expected as this originates from one single slice of data, whereas the other two planes originates from multiple slides of data, therefore making it harder to find a good fit.

The ANN was trained on the same 1D object described by Equation 15. It is highly likely that the in vivo data does not resemble this object in many ways, whereas the numerical brain phantom does. This could result in the network failing to find a good fit to the input data, leading to the images visible in Figure 18. This can probably be solved by finding a robust way of normalising the input data so that it more resembles the 1D training data. Further, it could also be possible to segment in vivo data of a healthy brain and use this as training data, and use other segments of the brain for validation. This could improve the accuracy of the super-resolution reconstruction for in vivo brain data.

It should also be noted that the architecture of the ANN was relatively simple, and could have probably been more optimised for this particular type of task. Based on the 1D-simulations we see that it should perform well both in terms of accuracy and precision, so a more delicate and optimised ANN for the particular task could still be a promising solution to the time problem mentioned previously from the gradient decent approach.

5.3 Conclusion and Future Aspects

With SERO we do not excite the slices serially, as is normally done. Instead, we shuffle them randomly resulting in the random overlap of slices giving us variable TR and intra-slice T_1 -weighting. In conclusion, SERO was found to have an improved accuracy in the lower SNR regions compared to conventional high-resolution direct sampling. For the diffusional variance map the accuracy of SERO was better than direct sampling across the scope of the simulations ($\text{SNR} \in [1, 12]$). However, the precision of direct sampling was better for the S_0 -, diffusivity- and variance maps. Furthermore, the inherent T_1 -weighting was also shown to help with the ill-posedness of the SRR problem in terms of finding a better solution to Equation 13. However, it also makes the fitting procedure more complicated due to its random nature and its increased dimensionality. In light of these results, SERO looks promising and could enable a fast method for DW-MRI that can achieve high-resolution maps of diffusivity. It exploits the improved SNR that is gained by sampling signal using thick slices. However, it suffers from a loss of precision compared to direct sampling due to the ill-posed SRR. Furthermore, due to the nature of the sampling matrix, the edges of the sampling space are undersampled.

SERO is a method in its infancy, and there is room for future development. During the scope of this study, a pulse sequence compatible with random slice positions and b-values was designed. However, the effectiveness of the pulse sequence was never thoroughly investigated. Thus, it could be of interest to keep designing the pulse sequence and try to optimise it further. It could also be of interest to keep investigating and optimise SERO further to increase the accuracy and precision in more complicated environments. Particularly, it could be of interest to keep working on a machine learning fitting approach to speed up the computational time. By designing an accurate machine learning based algorithm, one would vastly speed up the computational time of the SRR problem, while still maintaining a high precision.

References

- [1] Rebecca Smith-Bindman et al. “Trends in Use of Medical Imaging in US Health Care Systems and in Ontario, Canada, 2000-2016”. In: *JAMA* 322.9 (Sept. 2019), pp. 843–856. ISSN: 0098-7484. DOI: 10.1001/jama.2019.11456. URL: <https://doi.org/10.1001/jama.2019.11456>.
- [2] Denis Le Bihan. “Looking into the functional architecture of the brain with diffusion MRI”. en. In: *Nature Reviews Neuroscience* 4.6 (June 2003). Publisher: Nature Publishing Group, pp. 469–480. ISSN: 1471-0048. DOI: 10.1038/nrn1119. URL: <https://www.nature.com/articles/nrn1119>.
- [3] Roland Bammer. “Basic principles of diffusion-weighted imaging”. eng. In: *European Journal of Radiology* 45.3 (Mar. 2003), pp. 169–184. ISSN: 0720-048X. DOI: 10.1016/s0720-048x(02)00303-0.
- [4] Vinit Baliyan et al. “Diffusion weighted imaging: Technique and applications”. In: *World Journal of Radiology* 8.9 (Sept. 2016), pp. 785–798. ISSN: 1949-8470. DOI: 10.4329/wjr.v8.i9.785. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5039674/>.

-
- [5] Albert Macovski. “Noise in MRF”. en. In: *Magnetic Resonance in Medicine* 36.3 (1996). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910360327>, pp. 494–497. ISSN: 1522-2594. DOI: 10.1002/mrm.1910360327. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910360327>.
- [6] Hákon Gudbjartsson and Samuel Patz. “The Rician Distribution of Noisy MRI Data”. In: *Magnetic resonance in medicine* 34.6 (Dec. 1995), pp. 910–914. ISSN: 0740-3194. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2254141/>.
- [7] Esben Plenge et al. “Super-resolution methods in MRI: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time?” en. In: *Magnetic Resonance in Medicine* 68.6 (2012). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.24187>, pp. 1983–1993. ISSN: 1522-2594. DOI: 10.1002/mrm.24187. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24187>.
- [8] Gwendolyn Van Steenkiste et al. “Super-resolution reconstruction of diffusion parameters from diffusion-weighted images with different slice orientations”. eng. In: *Magnetic Resonance in Medicine* 75.1 (Jan. 2016), pp. 181–195. ISSN: 1522-2594. DOI: 10.1002/mrm.25597.
- [9] Samantha J. Holdsworth, Rafael O’Halloran, and Kawin Setsompop. “The quest for high spatial resolution diffusion-weighted imaging of the human brain in vivo”. eng. In: *NMR in biomedicine* 32.4 (Apr. 2019), e4056. ISSN: 1099-1492. DOI: 10.1002/nbm.4056.
- [10] H. Greenspan et al. “MRI inter-slice reconstruction using super-resolution”. In: *Magnetic Resonance Imaging* 20.5 (June 2002), pp. 437–446. ISSN: 0730-725X. DOI: 10.1016/S0730-725X(02)00511-8. URL: <https://www.sciencedirect.com/science/article/pii/S0730725X02005118>.
- [11] Kelvin J. Layton et al. “Pulseseq: A rapid and hardware-independent pulse sequence prototyping framework”. eng. In: *Magnetic Resonance in Medicine* 77.4 (Apr. 2017), pp. 1544–1552. ISSN: 1522-2594. DOI: 10.1002/mrm.26235.
- [12] CODATA Value: proton gyromagnetic ratio in MHz/T. URL: <https://physics.nist.gov/cgi-bin/cuu/Value?gammabar> (visited on 04/26/2024).
- [13] N. Bloembergen, E. M. Purcell, and R. V. Pound. “Relaxation Effects in Nuclear Magnetic Resonance Absorption”. In: *Physical Review* 73.7 (Apr. 1948). Publisher: American Physical Society, pp. 679–712. DOI: 10.1103/PhysRev.73.679. URL: <https://link.aps.org/doi/10.1103/PhysRev.73.679>.
- [14] S. H. Koenig et al. “Magnetic field dependence of $1/T_1$ of protons in tissue”. eng. In: *Investigative Radiology* 19.2 (1984), pp. 76–81. ISSN: 0020-9996. DOI: 10.1097/00004424-198403000-00002.
- [15] G. D. Fullerton, I. L. Cameron, and V. A. Ord. “Frequency dependence of magnetic resonance spin-lattice relaxation of protons in biological materials”. eng. In: *Radiology* 151.1 (Apr. 1984), pp. 135–138. ISSN: 0033-8419. DOI: 10.1148/radiology.151.1.6322223.
- [16] Jim Pipe. “Chapter 2 - Pulse Sequences for Diffusion-weighted MRI”. In: *Diffusion MRI*. Ed. by Heidi Johansen-Berg and Timothy E. J. Behrens. San Diego: Academic Press, Jan. 2009, pp. 11–35. ISBN: 978-0-12-374709-9. DOI: 10.1016/B978-0-12-374709-9.00002-X. URL: <https://www.sciencedirect.com/science/article/pii/B978012374709900002X>.
- [17] R. Turner, D. Le Bihan, and A. S. Chesnick. “Echo-planar imaging of diffusion and perfusion”. eng. In: *Magnetic Resonance in Medicine* 19.2 (June 1991), pp. 247–253. ISSN: 0740-3194. DOI: 10.1002/mrm.1910190210.
- [18] Martin H. Maurer and Johannes T. Heverhagen. “Diffusion weighted imaging of the prostate—principles, application, and advances”. In: *Translational Andrology and Urology* 6.3 (June 2017), pp. 490–498. ISSN: 2223-4691. DOI: 10.21037/tau.2017.05.06. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5503962/>.
-

-
- [19] Peter J. Basser and Derek K. Jones. “Diffusion-tensor MRI: theory, experimental design and data analysis - a technical review”. eng. In: *NMR in biomedicine* 15.7-8 (2002), pp. 456–467. ISSN: 0952-3480. DOI: 10.1002/nbm.783.
- [20] A. Einstein. “Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen”. en. In: *Annalen der Physik* 322.8 (1905). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19053220806>, pp. 549–560. ISSN: 1521-3889. DOI: 10.1002/andp.19053220806. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19053220806>.
- [21] Paul T. Callaghan. “Physics of Diffusion”. In: *Diffusion MRI: Theory, Methods, and Applications*. Ed. by PhD Jones Derek K. Oxford University Press, Nov. 2010, p. 0. ISBN: 978-0-19-536977-9. DOI: 10.1093/med/9780195369779.003.0004. URL: <https://doi.org/10.1093/med/9780195369779.003.0004>.
- [22] Denis Le Bihan and E. Breton. “Imagerie de diffusion in-vivo par résonance magnétique nucléaire”. In: *Comptes-Rendus de l’Académie des Sciences* 93.5 (Dec. 1985), pp. 27–34. URL: <https://hal.science/hal-00350090>.
- [23] D. Le Bihan et al. “MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders”. eng. In: *Radiology* 161.2 (Nov. 1986), pp. 401–407. ISSN: 0033-8419. DOI: 10.1148/radiology.161.2.3763909.
- [24] Jens H. Jensen et al. “Diffusional kurtosis imaging: The quantification of non-gaussian water diffusion by means of magnetic resonance imaging”. en. In: *Magnetic Resonance in Medicine* 53.6 (2005). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.20508>, pp. 1432–1440. ISSN: 1522-2594. DOI: 10.1002/mrm.20508. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.20508>.
- [25] Filip Szczepankiewicz et al. “The link between diffusion MRI and tumor heterogeneity: Mapping cell eccentricity and density by diffusional variance decomposition (DIVIDE)”. eng. In: *NeuroImage* 142 (Nov. 2016), pp. 522–532. ISSN: 1095-9572. DOI: 10.1016/j.neuroimage.2016.07.038.
- [26] Klaus Scheffler. “Superresolution in MRI?” eng. In: *Magnetic Resonance in Medicine* 48.2 (Aug. 2002), 408, author reply 409. ISSN: 0740-3194. DOI: 10.1002/mrm.10203.
- [27] Michele Nicastrò et al. “To shift or to rotate? Comparison of acquisition strategies for multi-slice super-resolution magnetic resonance imaging”. English. In: *Frontiers in Neuroscience* 16 (Nov. 2022). Publisher: Frontiers. ISSN: 1662-453X. DOI: 10.3389/fnins.2022.1044510. URL: <https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2022.1044510/full>.
- [28] W.H. Swann. “A survey of non-linear optimization techniques”. en. In: *FEBS Letters* 2.S1 (Mar. 1969). ISSN: 0014-5793, 1873-3468. DOI: 10.1016/0014-5793(69)80075-X. URL: <https://febs.onlinelibrary.wiley.com/doi/10.1016/0014-5793%2869%2980075-X>.
- [29] *Solve nonlinear curve-fitting (data-fitting) problems in least-squares sense - MATLAB lsqcurvefit - MathWorks Nordic*. URL: https://se.mathworks.com/help/optim/ug/lsqcurvefit.html#buuhcjo_seealso.
- [30] Kelvin J. Layton, Maxim Zaitsev, and Stefan Krobth. *pulseq/pulseq*. original-date: 2015-04-27T07:31:39Z. Apr. 2024. URL: <https://github.com/pulseq/pulseq>.
- [31] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin. “Adaptive reconstruction of phased array MR imagery”. eng. In: *Magnetic Resonance in Medicine* 43.5 (May 2000), pp. 682–690. ISSN: 0740-3194. DOI: 10.1002/(sici)1522-2594(200005)43:5<682::aid-mrm10>3.0.co;2-g.
- [32] François Chollet et al. *Keras*. <https://keras.io>. 2015.
-

-
- [33] Cristina Sainz Martinez, Meritxell Bach Cuadra, and João Jorge. “BigBrain-MR: a new digital phantom with anatomically-realistic magnetic resonance properties at 100- μm resolution for magnetic resonance methods development”. In: *NeuroImage* 273 (June 2023), p. 120074. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2023.120074. URL: <https://www.sciencedirect.com/science/article/pii/S1053811923002203>.
- [34] Geraline Vis et al. “Accuracy and precision in super-resolution MRI: Enabling spherical tensor diffusion encoding at ultra-high b-values and high resolution”. eng. In: *NeuroImage* 245 (Dec. 2021), p. 118673. ISSN: 1095-9572. DOI: 10.1016/j.neuroimage.2021.118673.
- [35] Nan-kuei Chen and Alice M. Wyrwicz. “Removal of EPI Nyquist ghost artifacts with two-dimensional phase correction”. eng. In: *Magnetic Resonance in Medicine* 51.6 (June 2004), pp. 1247–1253. ISSN: 0740-3194. DOI: 10.1002/mrm.20097.

A Appendix

Below are the high-resolution parameter maps acquired using SERO and an aspect ratio of 4.

